



ExtremeSwitching 200

Series: Administration Guide

Copyright © 2018 Extreme Networks, Inc. All rights reserved.

Legal Notice

Extreme Networks, Inc. reserves the right to make changes in specifications and other information contained in this document and its website without prior notice. The reader should in all cases consult representatives of Extreme Networks to determine whether any such changes have been made.

The hardware, firmware, software or any specifications described or referred to in this document are subject to change without notice.

Trademarks

Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries.

All other names (including any product names) mentioned in this document are the property of their respective owners and may be trademarks or registered trademarks of their respective companies/owners.

For additional information on Extreme Networks trademarks, please see:

www.extremenetworks.com/company/legal/trademarks

Software Licensing

Some software files have been licensed under certain open source or third-party licenses. End-user license agreements and open source declarations can be found at:

www.extremenetworks.com/support/policies/software-licensing

Support

For product support, phone the Global Technical Assistance Center (GTAC) at 1-800-998-2408 (toll-free in U.S. and Canada) or +1-408-579-2826. For the support phone number in other countries, visit: <http://www.extremenetworks.com/support/contact/>

For product documentation online, visit: <https://www.extremenetworks.com/documentation/>

Table of Contents

Preface.....	6
Conventions.....	6
Providing Feedback to Us.....	7
Getting Help.....	7
Extreme Networks Documentation.....	7
Chapter 1: Getting Started.....	9
Connecting the Switch to the Network.....	9
Booting the Switch.....	11
Understanding the Utility Menu.....	12
Understanding the User Interfaces.....	17
Chapter 2: Getting Started with Stacking.....	24
Understanding Switch Stacks.....	24
Switch Stack Software Compatibility Recommendations.....	26
Incompatible Software and Stack Member Image Upgrades.....	27
Switch Stack Configuration Files.....	27
Switch Stack Management Connectivity.....	27
General Practices.....	28
Initial Installation and Power-up of a Stack.....	28
Removing a Unit from the Stack.....	29
Adding a Unit to an Operating Stack.....	29
Replacing the Stack Member with a New Unit.....	29
Renumbering Stack Members.....	30
Moving a Manager to a Different Unit in the Stack.....	30
Removing a Manager Unit from an Operating Stack.....	31
Initiating a Warm Failover of the Manager Unit.....	31
Merging Two Operational Stacks.....	31
Preconfiguration.....	32
Chapter 3: Configuring System Information.....	33
Viewing the Dashboard.....	33
Viewing ARP Cache.....	35
Viewing Inventory Information.....	36
Viewing MAC Addresses.....	37
Viewing System Resources.....	37
Defining General Device Information.....	38
Managing the DHCP Server.....	73
Configuring DNS.....	81
Configuring Email Alerts.....	83
Configuring and Viewing ISDP Information.....	86
Configuring Link Dependency.....	89
Configuring Link Local Protocol Filtering.....	90
Configuring sFlow.....	92
Configuring SNTP Settings.....	97
Configuring Time Ranges.....	102
Configuring the Time Zone.....	104
Managing SNMP Traps.....	107

Managing CPU Traffic Filters.....	108
Viewing the System Firmware Status.....	112
Managing Logs.....	115
Configuring and Searching the Forwarding Database.....	121
Configuring Power Over Ethernet (PoE) and PoE Statistics.....	121
Viewing Device Port Information.....	125
Configuring and Viewing Device Slot Information.....	138
Defining SNMP Parameters.....	139
Viewing System Statistics.....	148
Using System Utilities.....	156
Chapter 4: Configuring Switching Information.....	167
Managing VLANs.....	167
Configuring UDLD.....	172
Private VLAN.....	174
Voice VLAN Configuration.....	177
Voice VLAN Interface.....	178
Port Auto Recovery.....	179
Creating MAC Filters.....	181
Configuring Dynamic ARP Inspection.....	182
GARP Configuration.....	187
Configuring DHCP Snooping.....	189
Configuring IPv6 DHCP Snooping.....	196
Configuring IGMP Snooping.....	201
Configuring IGMP Snooping Querier.....	205
Configuring MLD Snooping.....	208
Configuring MLD Snooping Querier.....	211
Creating Port Channels.....	214
Viewing Multicast Forwarding Database Information.....	218
Multicast VLAN Registration.....	220
Configuring Protected Ports.....	223
Configuring Spanning Tree Protocol.....	224
Mapping 802.1p Priority.....	236
Configuring Port Security.....	237
Managing LLDP.....	240
Loop Protection.....	250
Multiple Registration Protocol Configuration.....	251
Chapter 5: Configuring Routing.....	256
Configuring ARP.....	256
Configuring Global IP Settings.....	258
Router.....	267
Configuring Routing Information Protocol (RIP).....	271
Chapter 6: Managing Device Security.....	272
Port Access Control.....	272
RADIUS Settings.....	284
TACACS+ Settings.....	289
Authentication Manager.....	292
Chapter 7: Configuring IPv6.....	296
Global Configuration.....	296

Chapter 8: Configuring Quality of Service.....	298
Configuring Access Control Lists.....	298
Configuring Auto VoIP.....	311
Configuring Class of Service.....	315
Configuring Diffserv.....	317
Appendix A: Configuration Examples.....	331
Configuring VLANs.....	331
Configuring Multiple Spanning Tree Protocol.....	335
Configuring VLAN Routing.....	338
Configuring 802.1X Network Access Control.....	341
Configuring Authentication Tiering.....	342
Configuring Differentiated Services for VoIP.....	343
IGMP and MLD Snooping Switches.....	346
Configuring Port Mirroring.....	350
Bidirectional Forwarding Detection.....	351
Index.....	363
Glossary.....	353

Preface

This section discusses the conventions used in this guide, ways to provide feedback, additional help, and other Extreme Networks publications.

Conventions

This section discusses the conventions used in this guide.

Text Conventions

The following tables list text conventions that are used throughout this guide.

Table 1: Notice Icons

Icon	Notice Type	Alerts you to...
	General Notice	Helpful tips and notices for using the product.
	Note	Important features or instructions.
	Caution	Risk of personal injury, system damage, or loss of data.
	Warning	Risk of severe personal injury.
<i>New!</i>	New Content	Displayed next to new content. This is searchable text within the PDF.

Table 2: Text Conventions

Convention	Description
Screen displays	This typeface indicates command syntax, or represents information as it appears on the screen.
The words enter and type	When you see the word “enter” in this guide, you must type something, and then press the Return or Enter key. Do not press the Return or Enter key when an instruction simply says “type.”
[Key] names	Key names are written with brackets, such as [Return] or [Esc] . If you must press two or more keys simultaneously, the key names are linked with a plus sign (+). Example: Press [Ctrl]+[Alt]+[Del]
<i>Words in italicized type</i>	Italics emphasize a point or denote new terms at the place where they are defined in the text. Italics are also used when referring to publication titles.

Terminology

When features, functionality, or operation is specific to a switch family, such as ExtremeSwitching, the family name is used. Explanations about features and operations that are the same across all product families simply refer to the product as the *switch*.

Providing Feedback to Us

We are always striving to improve our documentation and help you work better, so we want to hear from you! We welcome all feedback but especially want to know about:

- Content errors or confusing or conflicting information.
- Ideas for improvements to our documentation so you can find the information you need faster.
- Broken links or usability issues.

If you would like to provide feedback to the Extreme Networks Information Development team about this document, please contact us using our short [online feedback form](#). You can also email us directly at documentation@extremenetworks.com.

Getting Help

If you require assistance, contact Extreme Networks using one of the following methods:

- **GTAC (Global Technical Assistance Center) for Immediate Support**
 - **Phone:** 1-800-998-2408 (toll-free in U.S. and Canada) or +1 408-579-2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact
 - **Email:** support@extremenetworks.com. To expedite your message, enter the product name or model number in the subject line.
- **Extreme Portal** — Search the GTAC knowledge base, manage support cases and service contracts, download software, and obtain product licensing, training, and certifications.
- **The Hub** — A forum for Extreme Networks customers to connect with one another, answer questions, and share ideas and feedback. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.

Before contacting Extreme Networks for technical support, have the following information ready:

- Your Extreme Networks service contract number and/or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any action(s) already taken to resolve the problem
- A description of your network environment (such as layout, cable type, other relevant environmental information)
- Network load at the time of trouble (if known)
- The device history (for example, if you have returned the device before, or if this is a recurring problem)
- Any related RMA (Return Material Authorization) numbers

Extreme Networks Documentation

To find Extreme Networks product guides, visit our documentation pages at:

Current Product Documentation	www.extremenetworks.com/documentation/
Archived Documentation (for earlier versions and legacy products)	www.extremenetworks.com/support/documentation-archives/
Release Notes	www.extremenetworks.com/support/release-notes

Open Source Declarations

Some software files have been licensed under certain open source licenses. More information is available at: www.extremenetworks.com/support/policies/software-licensing/.

1 Getting Started

Connecting the Switch to the Network
Booting the Switch
Understanding the Utility Menu
Understanding the User Interfaces

This section describes how to start the switch and access the user interface.

Connecting the Switch to the Network

To enable remote management of the switch through telnet, a web browser, or *SNMP (Simple Network Management Protocol)*, you must connect the switch to the network. The switch has no IP address by default, and *DHCP (Dynamic Host Configuration Protocol)* is disabled, so you must provide network information by connecting to the switch command-line interface (CLI) by using a local serial connection.

To access the switch over a network you must first configure it with network information (an IP address, subnet mask, and default gateway). You can configure the IP information using any of the following:

- BOOTP
- DHCP
- Terminal interface via the EIA-232 port

After you configure network information, such as the IP address and subnet mask, and the switch is physically and logically connected to the network, you can manage and monitor the switch remotely through SSH, telnet, a web browser, or an SNMP-based network management system. You can also continue to manage the switch through the terminal interface via the EIA-232 port.



Note

Some switches provide a dedicated service port for managing the switch. On switches without a dedicated service port, you use one of the network ports.

After you perform the physical hardware installation, you need to make a serial connection to the switch so that you can do one of the following:

- Manually configure network information for the management interface, or
- Enable the management interface as a DHCP or BootP client on your network (if not already enabled) and then view the network information after it is assigned by the DHCP server.

To connect to the switch and configure or view network information, use the following steps:

- 1 Using a straight-through modem cable, connect a VT100/ANSI terminal or a workstation to the console (serial) port.

If you attached a PC, Apple, or UNIX workstation, start a terminal-emulation program, such as HyperTerminal or TeraTerm.

2 Configure the terminal-emulation program to use the following settings:

- Baud rate: 9600 bps
- Data bits: 8
- Parity: none
- Stop bit: 1
- Flow control: none

3 Power on the switch.

For information about the boot process, including how to access the utility menu, see [Booting the Switch](#) on page 11.

4 Press **[Enter]**.

The `USER:` prompt displays.

Enter `admin` as the user name. There is no default password. Press **[Enter]** at the password prompt if you did not change the default password.

After a successful login, the screen shows the system prompt, for example `(Extreme 220) >`.

5 At the command prompt, enter `enable` to enter the Privileged EXEC command mode. There is no default password to enter Privileged EXEC mode. Press **[Enter]** at the password prompt if you did not change the default password.

The command prompt changes from `>` to `#`.

6 Configure network information.

If the unit has a service port:

- To have the address assigned through DHCP:

By default, the port is configured as a DHCP client. If your network has a DHCP server, then you need only to connect the switch to your network.

- To use BootP, change the protocol by entering: `serviceport protocol bootp`
- To disable DHCP/BootP and manually assign an IPv4 address, enter:

```
serviceport protocol none
serviceport ip ipaddress netmask [gateway]
```

For example:

```
serviceport ip 192.168.2.23 255.255.255.0 192.168.2.1
```

- To disable DHCP/BootP and manually assign an IPv6 address and (optionally) default gateway, enter:

```
serviceport protocol none
serviceport ipv6 address ipaddress/prefix-length [eui64]
serviceport ipv6 gateway gateway
```

- To view the assigned or configured network address, enter:

```
show serviceport
```

If the unit does not have a service port:

- To use a DHCP server to obtain the IP address, subnet mask, and default gateway information, enter:

```
network protocol dhcp
```

- To use a BootP server to obtain the IP address, subnet mask, and default gateway information, enter:

```
network protocol bootp
```

- To manually configure the IPv4 address, subnet mask, and (optionally) default gateway, enter:

```
network parms ipaddress netmask [gateway]
```

For example:

```
network parms 192.168.2.23 255.255.255.0 192.168.2.1
```

- To manually configure the IPv6 address, subnet mask, and (optionally) default gateway, enter:

```
network ipv6 address ipaddress/prefix-length [eui64]
network ipv6 gateway gateway
```

- To view the network information, enter `show network`
- To save these changes so that they are retained when the switch is rebooted, enter the following command:

```
copy system:running-config nvram:startup-config
```

After the switch is connected to the network, you can use the IP address for remote access to the switch by using a web browser or through Telnet or SSH.

Booting the Switch

When the power is turned on with the local terminal already connected, the switch goes through Power-On Self-Test (POST). POST runs every time the switch is initialized and checks hardware components to determine if the switch is fully operational before completely booting.

If a critical problem is detected, the program flow stops. If POST passes successfully, a valid executable image is loaded into RAM.

POST messages are displayed on the terminal and indicate test success or failure.

To boot the switch, perform the following steps:

- 1 Make sure that the serial cable is connected to the terminal.
- 2 Connect the power supply to the switch.
- 3 Power on the switch.

As the switch boots, the bootup test first counts the switch memory availability and then continues to boot.

- 4 During boot, you can use the Utility menu, if necessary, to run special procedures. To enter the Utility menu, press **2** within the first five seconds after the following message displays:

```
Select startup mode.  If no selection is made within 5 seconds,
the 200 Series Application will start automatically...
200 Series Startup -- Main Menu
1 - Start 200 Series Application
2 - Display Utility Menu
Select (1, 2): 2
```

For information about the Utility menu, see [Understanding the Utility Menu](#) on page 12.

- 5 If you do not start the Utility menu, the operational code continues to load.

After the switch boots successfully, the User login prompt appears and you can use the local terminal to begin configuring the switch. However, before configuring the switch, make sure that the software version installed on the switch is the latest version. If it is not the latest version, download and install the latest version. See [AutoInstall](#) on page 114.

Understanding the Utility Menu



Note

Utility menu functions vary on different operating systems and platforms. The following example might not represent the options available on your platform.

You can perform many configuration tasks through the Utility menu, which can be invoked after the first part of the POST is completed.

Use the following procedures to display the Utility menu:

- 1 During the boot process, press **2** within five seconds after the following message displays:

```
FASTPATH Startup
Select startup mode. If no selection is made within 5 seconds,
the FASTPATH Application will start automatically...
FASTPATH Startup -- Main Menu
1 - Start FASTPATH Application
2 - Display Utility Menu
Select (1, 2): 2
FASTPATH Startup
Options available
1 - Start FASTPATH Application
2 - Load Code Update Package
3 - Load Configuration
4 - Select Serial Speed
5 - Retrieve Error Log
6 - Erase Current Configuration
7 - Erase Permanent Storage
8 - Select Boot Method
9 - Activate Backup Image
10 - Start Diagnostic Application
11 - Reboot
12 - Erase All Configuration Files
Q - Quit from FASTPATH Startup
```

The following sections describe the Utility menu options. By default, if no selection is made within five seconds, the operational code starts.

Start the 200 Series Application

Use option 1 to resume loading the 200 Series application code.

To relaunch the boot process from the Utility menu, select **1** and press **[Enter]**.

The following prompt displays:

```
Extracting FASTPATH from image2.....done
Loading FASTPATH ...../mnt/application
done
Linking liblua.so to /lib/liblua.so
Linking libluaconn.so to /lib/libluaconn.so
```

```

Linking libproc_libs.so to /lib/libproc_libs.so
Linking librpcclt.so to /lib/librpcclt.so
Changing lighttpd file ownership to lighttpd:lighttpd
Expanding websrc.tar.gz into /mnt/www...done
PCI unit 0: Dev 0xb624, Rev 0x11, Chip BCM56624_B0, Driver BCM56624_B0
SOC unit 0 attached to PCI device BCM56624_B0
Adding BCM transport pointers
Configuring CPUTRANS TX
Configuring CPUTRANS RX

```

Load Code Update Package

Use option 2 when a new software version must be downloaded to replace corrupted files, update, or upgrade the system software.

To download software from the Utility menu:

- 1 On the **Utility** menu, select **2** and press **[Enter]**. The following prompt displays:

```
Select Mode of Transfer (Press T/X/Y/Z for TFTP/XMODEM/YMODEM/ZMODEM)
```

- 2 Select the transfer mode: press T/X/Y/Z for TFTP/XMODEM/YMODEM/ZMODEM.
- 3 When using HyperTerminal, click **Transfer** on the **HyperTerminal** menu bar.
- 4 From the **Transfer** menu, click **Send File**.

The **Send File** window opens.

- 5 Enter the file path for the file to be downloaded.
- 6 Make sure the protocol is defined per the transfer option selected in step 2 (XMODEM/YMODEM/ZMODEM).
- 7 Click **Send**.

The software is downloaded. Software downloading takes several minutes. The terminal emulation application, such as HyperTerminal, may display the loading process progress.

After software downloads, the switch reboots automatically.

Load Configuration

Use option 3 when a new configuration file must be downloaded to replace the saved system configuration file.

To download software from the Utility menu:

- 1 On the **Utility** menu, select **3** and press **[Enter]**.

The following prompt displays:

```
[Utility menu]4Ready to receive the file with XMODEM/CRC...
Ready to RECEIVE File tempcfg.bin in binary mode
Send several Control-X characters to cancel before transfer starts.
```

- 2 Select the transfer mode: press T/X/Y/Z for TFTP/XMODEM/YMODEM/ZMODEM.
- 3 When using HyperTerminal, click **Transfer** on the **HyperTerminal** menu bar.
- 4 From the **Transfer** menu, click **Send File**.

The **Send File** window opens.

- 5 Enter the file path for the file to be downloaded.
- 6 Make sure the protocol is defined per the transfer option selected in step 2 (XMODEM/YMODEM/ZMODEM).
- 7 Click **Send**.

The configuration file is downloaded. The terminal emulation application, such as HyperTerminal, may display the loading process progress.

Select Serial Speed

Use option **4** to change the baud rate of the serial interface.

To change the baud rate from the Utility menu:

- 1 On the **Utility menu**, select **4** and press **[Enter]**. The following prompt displays:

```
Select option (1-12 or Q):
1 - 2400
2 - 4800
3 - 9600
4 - 19200
5 - 38400
6 - 57600
7 - 115200
8 - Exit without change
```



Note

The selected baud rate takes effect immediately.

- 2 The bootup process resumes.

Retrieve Error Log

Use option 5 to retrieve the event log and download it to your ASCII terminal.

To retrieve the event log from the Utility menu:

- 1 On the **Utility menu**, select **5** and press **[Enter]**.

The following prompt displays:

```
Select Mode of Transfer (Press T/X/Y/Z for TFTP/XMODEM/YMODEM/ZMODEM)
```

- 2 Select the transfer mode (press T/X/Y/Z for TFTP/XMODEM/YMODEM/ZMODEM).

The following prompt displays:

```
File ascii.log.bin Ready to SEND in binary mode
Estimated File Size 169K, 1345 Sectors, 172032 Bytes
Estimated transmission time 3 minutes 20 seconds
Send several Control-X characters to cancel before transfer starts.
```

- 3 The bootup process resumes.

Erase Current Configuration

Use option 6 to load using the system default configuration and to boot without using the current startup configuration. Selecting 6 from the Utility menu restores system defaults. Boot Sequence can then be started by selecting 1 from the Utility menu.

To download software from the Utility menu:

- 1 On the **Utility** menu, select **6** and press **[Enter]**.

The following prompt displays:

```
Are you SURE you want to delete the configuration? (y/n):y
```

- 2 The bootup process resumes.

Erase Permanent Storage

Use option 7 to delete the active image from the flash memory. User action is confirmed with a Y/N question before executing the command.

To delete the backup image from the Utility menu:

- 1 On the **Utility** menu, select **7** and press **[Enter]**. The following prompt displays:

```
Are you SURE you want to delete operational code: image2? (y/n):y
Operational code deleted...
[Utility menu]
```

- 2 The bootup process resumes.

Select Boot Method

Use option 8 to select the method used to boot the system (FLASH, Network, or Serial boot). The default selection is FLASH.

To select the boot method from the Utility menu:

- 1 On the **Utility** menu, select **8** and press **[Enter]**. The following prompt displays:

```
Current boot method: FLASH
1 - Flash Boot
2 - Network Boot
3 - Serial Boot
4 - Exit without change
Select option (1-4):
```

- 2 The bootup process resumes.

Activate Backup Image

Use option 9 to activate the backup image. The active image becomes the backup when this option is selected.

To activate the backup image:

- 1 From the **Utility** menu, select **9** and press **[Enter]**. The following message displays:

```
Backup image - image2 activated.
```

- 2 The bootup process resumes.

Start Diagnostic Application

Use option 10 to run flash diagnostics. User action is confirmed with a Y/N question before executing the command.

To perform a complete test of the flash memory from the Utility menu:

- 1 On the **Utility** menu, select **10** and press **[Enter]**.

The following prompt displays:

```
Do you wish to run flash diagnostics? (Boot code region will not be tested.) (y/n):y
Input number of diagnostic iterations ->1
Testing 2 x 28F128J3 base: 0xfe000000
Iterations remaining = 1
Erasing sector 0
Verify sector 0 erased
Writing sector 0
Erasing sector 1
Verify sector 1 erased
Writing sector 1
Erasing sector 2
Verify sector 2 erased
Writing sector 2
Erasing sector 3
Verify sector 3 erased
Writing sector 3
Erasing sector 4
Verify sector 4 erased
Writing sector 4
Erasing sector 5
Verify sector 5 erased
Writing sector 5
Erasing sector 6
```

Note



This process runs until all sectors have been erased, verified erased, and written.

```
Flash Diagnostics passed
[Utility menu]
```

- 2 The bootup process resumes.

Reboot

Use option 11 to reboot the system:

- 1 From the **Utility** menu, select **11** and press **[Enter]**.
- 2 The bootup process resumes.

Erase All Configuration Files

Use option 12 to load using the system default configuration and to boot without using the current startup configuration. Selecting 12 from the Utility menu restores system defaults. Boot Sequence can then be started by selecting 1 from the Utility menu.

To download software from the Utility menu:

- 1 On the **Utility menu**, select **12** and press **[Enter]**.

The following prompt displays:

```
Are you SURE you want to delete the configuration? (y/n):y
```

- 2 The bootup process resumes.

Understanding the User Interfaces

The software includes a set of comprehensive management functions for configuring and monitoring the system by using one of the following methods:

- Web User Interface
- Command-Line Interface (CLI)
- SNMP
- OpEN API
- RESTful API
- RESTCONF
- NETCONF

Each of the standards-based management methods allows you to configure and monitor the components of the software. The method you use to manage the system depends on your network size and requirements, and on your preference.



Note

Not all components can be managed by each interface.

This guide describes how to use the web-based interface to manage and monitor the system. For information about how to manage and monitor the system by using the CLI, see *ExtremeSwitching 200 Series: Command Reference Guide*.



Note

The web configuration and monitoring pages and CLI commands available for each platform depend on the software version and modules installed. For more information about the modules, see *Getting Started* on page 9.

Using the Web Interface

To access the switch by using a web browser, the browser must meet the following software requirements:

- HTML version 4.0, or later

- HTTP version 1.1, or later
- JavaScript version 1.5, or later

Use the following procedures to log on to the web Interface:

- 1 Open a web browser and enter the IP address of the switch in the web browser address field.
- 2 Type the user name and password into the fields on the login screen, and then click **Login**.
The user name and password are the same as those you use to log on to the command-line interface. By default, the user name is `admin`, and there is no password. Passwords are case sensitive.
- 3 After the system authenticates you, the **System Description** window opens.

Figure 1 shows the layout of the web interface. Each web page contains three main areas: device view, the navigation menu, and the configuration status and options.

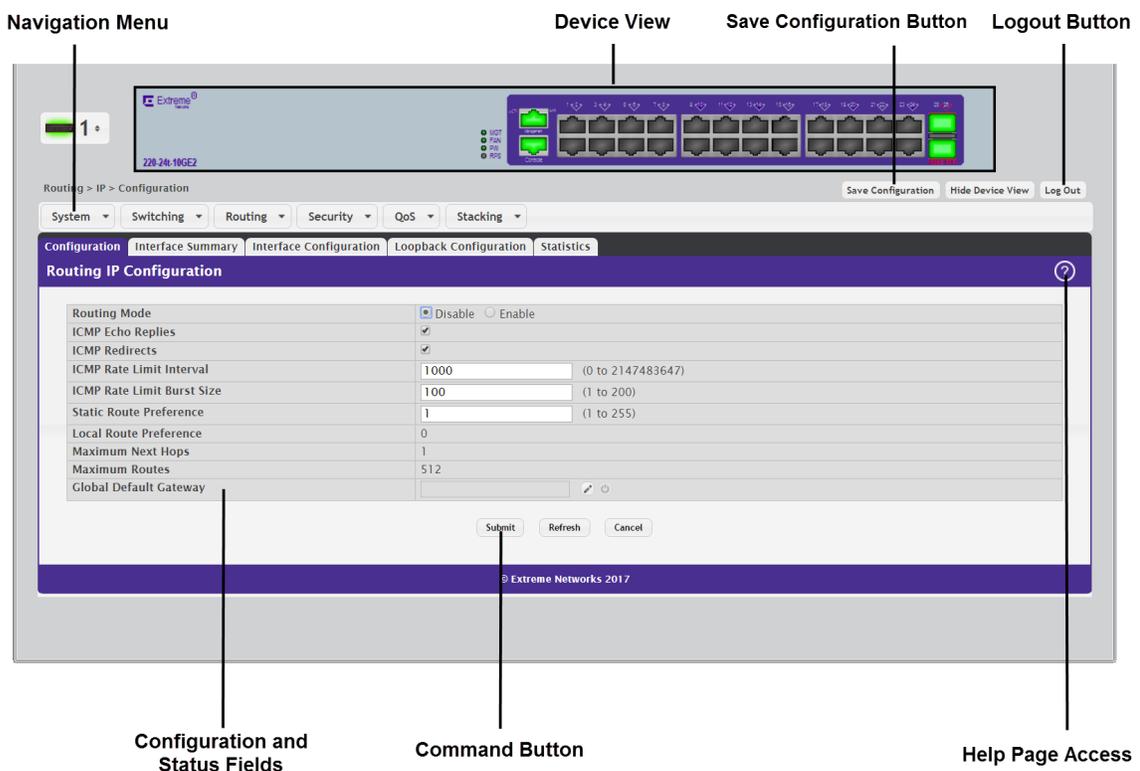


Figure 1: Web Interface Layout

Device View

The Device View is a Java applet that displays the ports on the switch. This graphic appears at the top of each page to provide an alternate way to navigate to port related configuration and monitoring options. The graphic also provides information about device ports, current configuration and status, table information, and feature components.

The port coloring indicates if a port is currently active. Green indicates that the port is enabled, red indicates that an error has occurred on the port, and blue indicates that the link is disabled.

Click the port you want to view or configure to see a menu that displays statistics and configuration options. Click the menu option to access the page that contains the configuration or monitoring options.

If you click the graphic but do not click a specific port, the main menu opens, as [Figure 2](#) shows. This menu contains the same option as the navigation menu on the left side of the page.

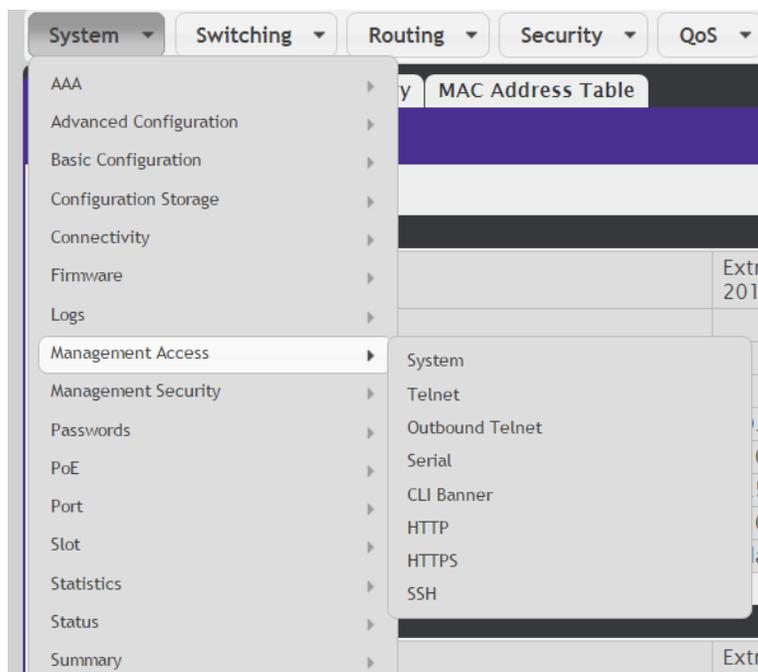


Figure 2: Management Access Menu

Navigation Menu

The navigation menu is on the top of the web interface. The navigation menu contains a list of various device features. The main items in the navigation menu can be expanded to view all the components under a specific feature, or retracted to hide the feature's components.

The navigation menu consists of a combination of main feature menus, submenus, and configuration and status pages. Click the feature menu, such as System or Switching, to view the options in that menu. Each menu contains submenus, HTML pages, or a combination of both. [Figure 3](#) shows an example of a feature menu (Switching), submenu (VLAN), and the active page in the navigation menu (Port Configuration).

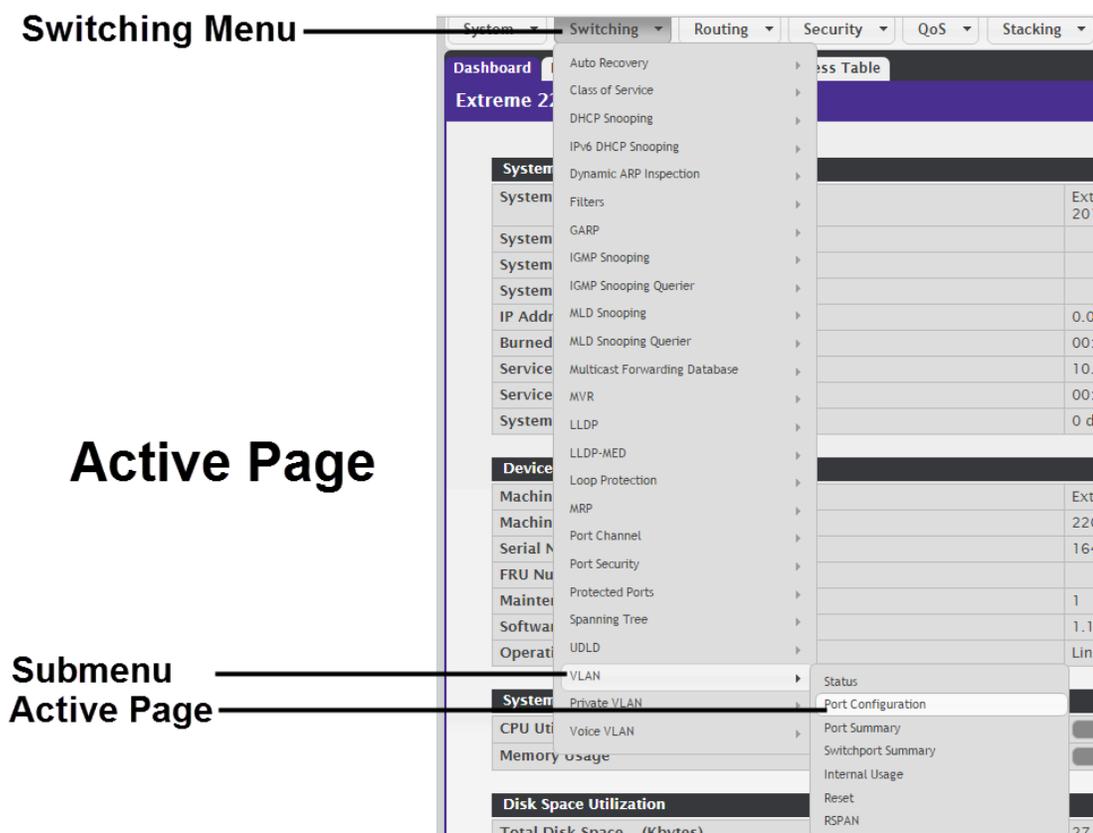


Figure 3: Navigation Menu View

When you click a menu or submenu, the color turns from gray to red, the menu expands to show its contents, and the arrow on the right side of the menu rotates. If you click a page under a menu or submenu, a new window opens in the main frame.

Configuration and Status Fields

The main area of the screen displays the fields you use to configure or monitor the switch. On pages that contain configuration options, you can input information into fields or select options from drop-down menus.

Each page contains access to the HTML-based help that explains the fields to configure or view on the page. Many pages also contain command buttons.

Table 3 shows the command buttons that are used throughout the pages in the web interface.

Table 3: Common Command Buttons

Button	Function
Submit	Sends the updated configuration to the switch. Configuration changes take effect immediately, but changes are not retained across a power cycle unless you save them to the system configuration file. To save the configuration to non-volatile memory, click Save Configuration in the upper-right portion of any configuration window.
Refresh	Refreshes the page with the most current information.
Delete	Removes the selected entry from the running configuration.
Clear	Removes all entries from a table or resets statistical counters to the default value.
Edit	Changes an existing entry.
Remove	Deletes the selected entries.
Clear Counter	Clears all the statistics counters, resetting all switch summary and detailed statistics to default values.
Logout	Ends the session.

**Caution**

Submitting changes makes them effective during the current boot session only. You must save any changes if you want them to be retained across a power cycle (reboot).

Table Sorting

Tables shown in the web pages can be sorted in each column. To sort a column, click at the top of the column to sort by that field. For example, in the **Event Log** page, clicking on the Event Time column will sort the entries by that field.

Help Page Access

The Help button, shown in [Figure 4](#), is always available in the upper-right corner of the active page. Click the button to open a new page that contains information about the configuration fields, status fields, and command buttons available on the active page.

The online help pages are context sensitive. For example, if the **IP Addressing** page is open, the help topic for that page displays when you click the Help button.

**Figure 4: Help Button**

[Figure 1](#) on page 18 shows the location of the Help link on the web interface.

User-Defined Fields

User-defined field names can contain between 1 and 159 characters, unless otherwise noted on the configuration web page.

All characters may be used except the following (unless specifically noted for a particular feature):

- \
- /
- *
- ?
- <
- >
- |

Using SNMP

For 200 Series software that includes the [SNMP](#) module, you can configure SNMP groups and users that can manage traps that the SNMP agent generates.

200 Series uses both standard public MIBs for standard functionality and private MIBs that support additional switch functionality. All private MIBs begin with a "-" prefix. The main object for interface configuration is in -SWITCHING-MIB, which is a private MIB. Some interface configurations also involve objects in the public MIB, IF-MIB.

SNMP is enabled by default. The System Description web page, which is the page the displays after a successful login and the `show sysinfo` command display the information you need to configure an SNMP manager to access the switch.

Any user can connect to the switch using the SNMPv3 protocol, but for authentication and encryption, you need to configure a new user profile. To configure a profile by using the CLI, see the SNMP section in [ExtremeSwitching 200 Series: Command Reference Guide](#).

To configure an SNMPv3 profile by using the web interface, use the following steps:

- 1 Select **System > Users > Accounts** from the navigation menu on the left side of the web interface.
- 2 From the **Accounts** menu, select **Add** to create a new user.
- 3 Enter a new user name in the **User Name** field.
- 4 Enter a new user password in the **Password** field and then retype it in the **Confirm** field.
To use SNMPv3 Authentication for this user, set a password of eight or more alphanumeric characters.
- 5 To enable authentication, use the **Authentication Protocol** menu to select either [MD5 \(Message-Digest algorithm 5\)](#) or SHA for the authentication protocol.
- 6 To enable encryption, use the **Encryption Protocol** menu to select **DES** for the encryption scheme. Then, enter an encryption code of eight or more alphanumeric characters in the **Encryption Key** field.
- 7 Click **Submit**.

To access configuration information for SNMPv1 or SNMPv2, click and click the page that contains the information to configure.

Using the Command-Line Interface

The CLI is a text-based way to manage and monitor the system. You can access the CLI by using a direct serial connection or by using a remote logical connection with Telnet or SSH.

The CLI groups commands into modes according to the command function. Each of the command modes supports specific software commands. The commands in one mode are not available until you switch to that particular mode, with the exception of the User EXEC mode commands. You can execute the User EXEC mode commands in the Privileged EXEC mode.

To display the commands available in the current mode, enter a question mark (?) at the command prompt. To display the available command keywords or parameters, enter a question mark (?) after each word you type at the command prompt. If there are no additional command keywords or parameters, or if additional parameters are optional, the following message appears in the output:

```
<cr>                               Press Enter to execute the command
```

For more information about the CLI, see [ExtremeSwitching 200 Series: Command Reference Guide](#). That guide lists each available command with the following information:

- The command keywords and the required and optional parameters.
- The command mode you must be in to access the command.
- The default value, if any, of a configurable setting on the device.

The `show` commands in the guide also include a description of the information displayed by the command.

2 Getting Started with Stacking

Understanding Switch Stacks

Switch Stack Software Compatibility Recommendations

Incompatible Software and Stack Member Image Upgrades

Switch Stack Configuration Files

Switch Stack Management Connectivity

General Practices

Initial Installation and Power-up of a Stack

Removing a Unit from the Stack

Adding a Unit to an Operating Stack

Replacing the Stack Member with a New Unit

Renumbering Stack Members

Moving a Manager to a Different Unit in the Stack

Removing a Manager Unit from an Operating Stack

Initiating a Warm Failover of the Manager Unit

Merging Two Operational Stacks

Preconfiguration

This section describes the concepts and recommended operating procedures to manage stacked Ethernet switches running 200 Series.



Note

For complete syntax and usage information for the commands used in this chapter, see [ExtremeSwitching 200 Series: Command Reference Guide](#).

Understanding Switch Stacks

A switch stack is a set of up to four Ethernet switches connected through their stacking ports. One of the switches controls the operation of the stack and is called the stack manager. All other switches in the stack are stack members. The stack members use stacking technology to behave and work together as a unified system. Layer 2 and Layer 3 protocols present the entire switch stack as a single entity to the network.

The stack manager is the single point of stack-wide management. From the stack manager, you configure:

- System-level (global) features that apply to all stack members
- Interface-level features for all interfaces on any stack member

A switch stack is identified in the network by its network IP address. The network IP address is assigned according to the MAC address of the stack manager. The MAC address used by the switch is the MAC

address of the manager. You can see this address by issuing the `show network` command. Every stack member is uniquely identified by its own stack member number.

All stack members are eligible stack managers. Exception: Setting a stack member's priority to 0 (zero) makes it ineligible for manager selection. When the stack is formed, one of the units is automatically selected as the standby for the stack. The standby of the stack takes over as manager if the current manager fails. The standby of the stack can also be configured using the command `standby unit-number`.

The stack manager contains the saved and running configuration files for the switch stack. The configuration files include the system-level settings for the switch stack and the interface-level settings for all stack members. Each stack member retains a copy of the saved file for backup purposes.

If the manager is removed from the stack, the standby of the stack will take over and will then run from that saved configuration.

You can use these methods to manage switch stacks:

- Web interface
- CLI over a serial connection to the console port of the manager
- A network management application through the *SNMP (Simple Network Management Protocol)*

Switch Stack Membership

A switch stack has up to n stack members, including the manager, connected through their stacking ports. A switch stack always has one stack manager.

A standalone switch is a switch stack with one stack member that also operates as the stack manager. You can connect one standalone switch to another to create a switch stack containing two stack members, with one of them being the stack manager. You can connect standalone switches to an existing switch stack to increase the stack membership.

If you replace a stack member with an identical model, the new switch functions with exactly the same configuration as the replaced switch, assuming that the new switch is using the same member number as the replaced switch. By default, 200 Series configures the new member.

The operation of the switch stack continues uninterrupted during membership changes unless you remove the stack manager.

Stack Manager Election and Re-Election

The stack manager is elected or re-elected based on one of these factors and in the order listed:

- The switch that is currently the stack manager.
- The switch with the highest stack member priority value.



Note

We recommend assigning the highest priority value to the switch that you prefer to be the stack manager. This ensures that the switch is re-elected as stack manager if a re-election occurs.

- The switch with the higher MAC address.

A stack manager retains its role unless one of these events occurs:

- The stack manager is removed from the switch stack.
- The stack manager is rebooted or powered off.
- The stack manager has failed.
- The switch stack membership is increased by adding powered-on standalone switches or switch stacks.

In the case of a manager re-election, the new stack manager becomes available after a few seconds.

If a new stack manager is elected and the previous stack manager becomes available, the previous stack manager does not resume its role as stack manager.

Stack Member Numbers

A stack member number (1 to n) identifies each member in the switch stack. The member number also determines the interface-level configuration that a stack member uses. You can display the stack member number by using the `show switch` Privileged EXEC command.

A new, out-of-the-box switch (one that has not joined a switch stack or has not been manually assigned a stack member number) ships with a default stack member number of 1. When it joins a switch stack, its default stack member number changes to the lowest available member number in the stack.

Stack members in the same switch stack cannot have the same stack member number. Every stack member, including a standalone switch, retains its member number until you manually change the number or unless the number is already being used by another member in the stack. See [Renumbering Stack Members](#) on page 30 and [Merging Two Operational Stacks](#) on page 31.

Stack Member Priority Values

You can set the stack member's priority in the range 0 to 15.



Note

Setting the switch priority to 0 (zero) makes it ineligible for manager selection.

Switch Stack Software Compatibility Recommendations

All stack members must run the same 200 Series software version to ensure compatibility between stack members. The software versions on all stack members, including the stack manager, must be the same. This helps ensure full compatibility in the stack protocol version among the stack members.

If a stack member is running a software version that is not the same as the stack manager, then the stack member joins the stack but stays in code incompatible status (the stack unit is not allowed to join the stack as a fully functional member). Use the `show switch` command to list the stack members and the software versions. The new unit will be visible. The administrator can load the code to that new unit and reboot the unit. The ports on the unit in software mismatch state do not come up.

Incompatible Software and Stack Member Image Upgrades

You can upgrade a switch that has an incompatible software image by using the `copy {active | backup} unit://unit-number/{active | backup}` command from config stack mode. It copies the software image from an existing stack member to the one with incompatible software. Because that switch does not automatically reload, issue a reload command to that switch and it joins the stack as a fully functioning member.

Switch Stack Configuration Files

The configuration files record settings for all global and interface specific settings that define the operation of the stack and individual members. Once a save to the configuration is issued, all stack members store a copy of the configuration settings. If a stack manager becomes unavailable, any stack member assuming the role of stack manager will operate from the saved configuration files.

When a new, out-of-box switch joins a switch stack, it uses the system-level settings of that switch stack. For the switch to store this system-level configuration, you must issue the following command:

```
copy system:running-config nvram:startup-config
(in Privileged EXEC)
```

This command saves passwords and all other changes to the device.

If you do not save the configuration using this command, all configurations will be lost when a power cycle is performed on the networking device or when the networking device is rebooted.

Note



After downloading a configuration file to a stack, you must perform a configuration save operation from the 200 Series user interface (that is, the `copy` command shown above) to distribute this configuration to non-management units in the stack. This is also true of SSH key files and SSL certificate files. From the command line interface, the following command can be used:

```
copy system:running-config nvram:startup-config (in Privileged EXEC)
```

You back up and restore the stack configuration in the same way as you would for standalone switch configuration.

Switch Stack Management Connectivity

You manage the switch stack and the stack member interfaces through the stack manager. You can use the web interface, CLI, and *SNMP*. You cannot manage stack members on an individual switch basis.

Connectivity to the Switch Stack through Console Ports

You can connect to the stack manager through the console port of the stack manager only.

Connectivity to the Switch Stack through Telnet

You can also Telnet to the stack manager using the commands `telnet ipaddress` and `login`.

General Practices

The following practices are highly recommended:

- When issuing a command (such as move management, or renumber), we recommend that you allow the command to fully complete before issuing the next command. For example, if you reboot a stack member, verify - using the `show port` command - that the switch has remerged with the stack and all ports are joined before issuing the next command.
- When physically removing or relocating a unit, always power down the unit before disconnecting stack cables.
- When reconnecting stack cables, connect them before powering up the unit, if possible. Tighten all connector screws, where applicable, to ensure a good connection.

The following sections provide switch stack configuration scenarios. Most of the scenarios assume at least two switches are connected through their stacking ports.

Initial Installation and Power-up of a Stack

Use the following steps to install and power-up a stack of switches:

- 1 Install units in rack whenever possible to prevent the units and cables from being disturbed
- 2 Install all stacking cables. Fully connect all cables, including the redundant stack link. We highly recommend that a redundant link be installed because this provides stack resiliency.
- 3 Identify the unit to be the manager. Power this unit up first.
- 4 To set up a stack, complete the following steps:
 - a Make sure there is a 200 Series image on each box.
 - b If the image does not exist or needs to be updated, use TFTP or xmodem to perform the update operation.
- 5 Monitor the console port. Allow this unit to come up to the login prompt. If the unit has the default configuration, it should come up as unit #1, and will automatically become a manager unit. If not, renumber the unit as desired.
- 6 If desired, preconfigure other units to be added to the stack. See [Preconfiguration](#) on page 32.
- 7 Power on a second unit, making sure it is adjacent (next physical unit in the stack) to the unit already powered up. This will ensure the second unit comes up as a member of the stack, and not a “Manager” of a separate stack.
- 8 Monitor the manager unit to see that the second unit joins the stack. Use the `show switch` command to determine when the unit joins the stack. It will be assigned a unit number (unit #2, if it has the default configuration).
- 9 If desired, renumber this stack unit. See [Renumbering Stack Members](#) on page 30 for recommendations for renumbering stack members.
- 10 Repeat steps 6 through 9 to add additional members to the stack. Always power on a unit adjacent to the units already in the stack.

Removing a Unit from the Stack

Use the following steps to remove a switch from the stack:

- 1 Make sure the redundant stack connection is in place and functional. All stack members should be connected in a logical ring.
- 2 Power down the unit to be removed.
- 3 Disconnect the stacking cables
- 4 If the unit is not to be replaced, reconnect the stack cable from the stack member above to the stack member below the unit being removed.
- 5 Remove the unit from the rack.
- 6 If desired, remove the unit from the configuration by issuing the command `no member unit-id` in Stack mode.

Using the web Interface, you can remove a member of the stack by selecting **Stacking > Base > Summary**. Then check the box next to the switch and click **Remove**.

Adding a Unit to an Operating Stack

Use the following steps to add a switch to a stack of switches while the stack is running:

- 1 Make sure that the redundant stack link is in place and functional. All stack members should be connected in a logical ring.
- 2 Preconfigure the new unit, if desired.
- 3 Install the new unit in the rack (assumes installation below the bottom-most unit, or above the top-most unit).
- 4 Disconnect the redundant stack cable that connects the last unit in the stack back up to the first unit in the stack at the new position in the ring where the new unit is to be inserted.
- 5 Connect this cable to the new unit, following the established order of connections. In other words, use the redundant stack cable to connect from the first box in the stack to the last.
- 6 Power up the new unit. Verify, by monitoring the manager unit console port, that the new unit successfully joins the stack by using the `show switch` command in EXEC mode. The new unit should always join as a "member" (never as manager; the existing manager of the stack should not change).
- 7 If the 200 Series software version of the newly added member is not the same as the existing stack, update the software image.

Adding a powered-up standalone unit to an operational stack is similar to merging two operational stacks where the standalone unit is a stack of one unit. See [Merging Two Operational Stacks](#) on page 31 for more details.

Using the web Interface, you can create a new member for the stack by selecting **Stacking > Base > Summary**. Then click **Add**, supply the attributes for the new member, and click **Submit**.

Replacing the Stack Member with a New Unit

There are two options here. If a stack member of a certain model number is replaced with another unit of the same model, follow these steps:

- 1 Follow the process in [Removing a Unit from the Stack](#) on page 29 to remove the desired stack member.
- 2 Follow the process in [Adding a Unit to an Operating Stack](#) on page 29 to add a new member to the stack with the following exceptions:
 - Insert the new member in the same position in the stack as the one removed.
 - It is not necessary to preconfigure the unit.

If a stack member is replaced with a unit of a different model number, follow these steps:

- 3 Follow the process in [Removing a Unit from the Stack](#) on page 29 to remove the desired stack member.
- 4 Remove the now-absent stack member from the configuration by issuing the `no member` command in Config Stack mode.
- 5 Add the new stack unit to the stack using the process described in [Adding a Unit to an Operating Stack](#) on page 29. The unit can be inserted into the same position as the unit just removed, or the unit can be inserted at the bottom of the stack. In either case, make sure all stack cables are connected with the exception of the cable at the position where the new unit is to be inserted to ensure that the stack does not get divided into two separate stacks, causing the election of a new manager.

Renumbering Stack Members

- 1 If particular numbering is required, we recommend that stack members be assigned specific numbers when they are first installed and configured in the stack, if possible.
- 2 If the desired stack unit number for a particular unit is unused, a unit can be renumbered simply by using the command `switch oldunit-id renumber newunit-id` in Global Config mode.
- 3 Renumbering a non-manager unit requires the unit to be rebooted for the renumbering to take effect. Renumbering a manager unit requires a reboot of all the switches in the stack for the renumbering to take effect.
- 4 If the newunit-id has been preconfigured, you may need to remove the newunit-id from the configuration before renumbering the unit.
- 5 If reassignment of multiple existing stack unit numbers is necessary, there are a number of implications in terms of mismatching of configuration. In this case, we recommend that all units except the manager be powered down and added back one at a time using the procedure in [Adding a Unit to an Operating Stack](#) on page 29.

Using the web Interface, you can renumber a switch using the **Stacking - Unit Configuration** page. To renumber a switch:

- 6 Select **Stacking > Base > Unit Configuration**.
- 7 In the **Switch ID** field, select the switch you want to renumber.
- 8 Click the Edit Switch ID (pencil) icon, select a new number in the drop-down list, and click **Submit**.

Moving a Manager to a Different Unit in the Stack

Use the following steps to change the stack manager from the current switch to a new switch in the stack:

- 1 Using the `movemanagement` command, move the manager to the desired unit number. The operation may take three minutes or longer depending on the stack size and configuration. The command, in Config Stack mode, is `movemanagement fromunit-id tounit-id`.
- 2 Make sure that you can log in on the console attached to the new manager. Use the `show switch` command to verify that all units rejoined the stack.
- 3 We recommend that you reboot the stack with the `reload` command in Privileged EXEC mode after moving the manager.

Removing a Manager Unit from an Operating Stack

Use the following steps to remove the manager unit from the stack during operation:

- 1 Move the designated manager to a different unit in the stack by following the steps in [Moving a Manager to a Different Unit in the Stack](#) on page 30.
- 2 Remove the unit from the stack by following the steps in [Removing a Unit from the Stack](#) on page 29.

Initiating a Warm Failover of the Manager Unit

You can use the `initiate failover` command to initiate a "warm" restart. This command reloads the management unit, triggering the standby unit to take over. As the standby management unit takes over, the system continues to forward end-user traffic. The end-user data streams may lose a few packets during the failure, but they do not lose their IP sessions, such as VoIP calls.

If there no standby unit is available when the initiate failover command is issued, the command fails with an error message stating that no standby unit exists. If the standby unit is not ready for a warm restart, the command fails with a similar error message. The `move management` command triggers a cold restart, even if the target unit is the backup unit.

Merging Two Operational Stacks

We recommend using the following procedure for merging two operational stacks:

- 1 Always power off all units in one stack before connecting to another stack.
- 2 Add the units as a group by unplugging one stacking cable in the operational stack and physically connecting all unpowered units.
- 3 Completely cable the stacking connections, making sure the redundant link is also in place.

Two operational stacks can also be merged by reconnecting stack cables without powering down all units in one stack. Connecting a powered-up standalone unit to an existing stack leads to same behavior as when merging two operational stacks. In such cases, the Manager re-election is done based on the rules listed in [Stack Manager Election and Re-Election](#) on page 25. One of the two managers wins the election and the losing stack manager reboots itself and all its member units. After the reboot, all of the losing stack members join the winning stack to form a single stack. The winning stack remains functional through the merge process. If the stack merge is performed in this way, then we strongly recommend setting the priority of the desired winner stack manager to a higher value than the stack manager that should lose the election.

Preconfiguration

This section explains how to configure units. Units do not necessarily have to be preconfigured in order to be added to the stack.

General information: All configurations on the stack, except unit numbers, are stored on the management unit. This means that a stack unit may be replaced with another device of the same type without having to reconfigure the switch. Unit numbers are stored independently on each switch, so that after power cycling the stack, the units always come back with the same unit numbers. The unit type associated with each unit number may be learned by the management unit automatically as the units are connected or preconfigured by the administrator.

- 1 Issue the command `member unit-id switchindex` to preconfigure a unit from the config stack mode. Supported unit types are shown by the `show supported switchtype` command.
 - To display supported switches, use Privileged EXEC mode. Enter the command `show supported switchtype SID`.
 - To add a new member (see [Adding a Unit to an Operating Stack](#) on page 29), use Config stack mode. Enter the command `member unit-id`.
- 2 Next, configure the unit you just defined with configuration commands, just as if the unit were physically present.
- 3 Ports for the preconfigured unit come up in "detached" state and can be seen with the `show port all` command in Privileged EXEC mode. The detached ports may now be configured for VLAN membership and any other port-specific configuration.
- 4 After a unit type is preconfigured for a specific unit number, attaching a unit with a different unit type for this unit number causes the switch to report an error. The Privileged Exec mode `show switch` command indicates "config mismatch" for the new unit and the ports on that unit do not come up. To resolve this situation, you may change the unit number of the mismatched unit, using the procedure in [Renumbering Stack Members](#) on page 30, or delete the preconfigured unit type using the command `no member unit-id` from the config stack mode.

3 Configuring System Information

Viewing the Dashboard
Viewing ARP Cache
Viewing Inventory Information
Viewing MAC Addresses
Viewing System Resources
Defining General Device Information
Managing the DHCP Server
Configuring DNS
Configuring Email Alerts
Configuring and Viewing ISDP Information
Configuring Link Dependency
Configuring Link Local Protocol Filtering
Configuring sFlow
Configuring SNTP Settings
Configuring Time Ranges
Configuring the Time Zone
Managing SNMP Traps
Managing CPU Traffic Filters
Viewing the System Firmware Status
Managing Logs
Configuring and Searching the Forwarding Database
Configuring Power Over Ethernet (PoE) and PoE Statistics
Viewing Device Port Information
Configuring and Viewing Device Slot Information
Defining SNMP Parameters
Viewing System Statistics
Using System Utilities

Use the features in the System menu to define the switch's relationship to its environment.

Viewing the Dashboard

After a successful login, the **Dashboard** window opens. This page provides a brief overview of the system.

To navigate to this page, click **System > Summary > Dashboard** in the navigation menu.

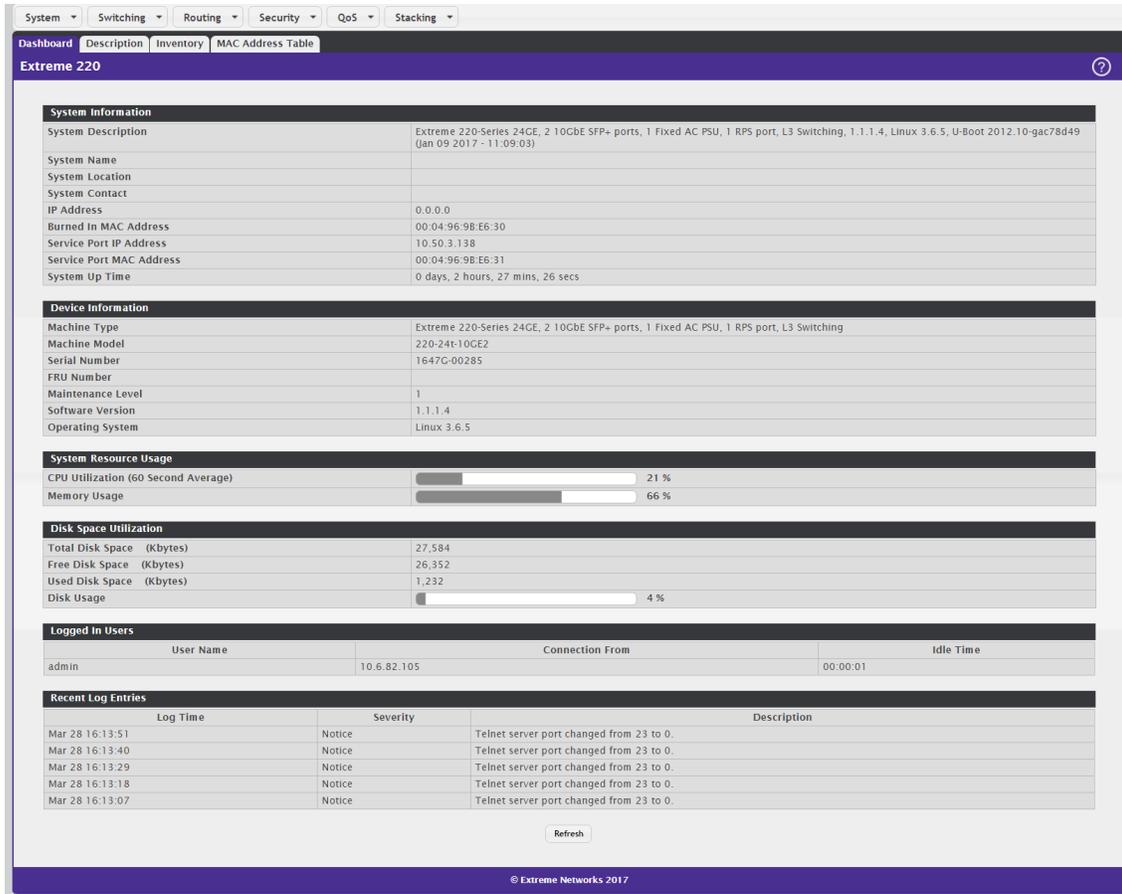


Figure 5: System Dashboard

Table 4: Dashboard Fields

Field	Description
System Information	
System Description	The product name of this device.
System Name	The configured name used to identify this device.
System Location	The configured location of this device.
System Contact	The configured contact person for this device.
IP Address	The IP address assigned to the network interface. The network interface is the logical interface that allows remote management of the device via any of the front-panel switch ports.
Burned In MAC Address	The device burned-in universally-administered media access control (MAC) address of the base system.
Service Port IP Address	The IP address assigned to the service port. The service port provides remote management access to the device. Traffic on this port is segregated from operational network traffic on the switch ports and cannot be switched or routed to the operational network.
Service Port MAC Address	The device burned-in universally-administered media access control (MAC) address of the service port.

Table 4: Dashboard Fields (continued)

Field	Description
System Up Time	The time in days, hours, minutes and seconds since the system was last rebooted.
Device Information	
Machine Type	The device hardware type or product family.
Machine Model	The model identifier, which is usually related to the Machine Type.
Serial Number	The unique device serial number.
FRU Number	The field replaceable unit number.
Maintenance Level	The device hardware change level identifier.
Software Version	The release.version.maintenance number of the software currently running on the device. For example, if the release is 1, the version is 2 and the maintenance number is 4, this version number is displayed as 1.2.4.
Operating System	The device operating system type and version identification information.
System Resource Usage	
CPU Utilization (60 Second Average)	The percentage of CPU utilization for the entire system averaged over the past 60 seconds.
Memory Usage	The percentage of total available system memory (RAM) that is currently in use.
Disk Space Utilization	
Disk Usage	The percentage of total available disk space that is currently in use.
Additional Fields	
Logged In Users	A brief summary indicating all other users currently logged into the device. The Idle Time field gives an indication of user activity, with a smaller time value denoting more recent access to the system.
Recent Log Entries	A brief list of the newest entries recorded in the system log.

Click **Refresh** to reload the page and refresh the Dashboard.

Viewing ARP Cache

The *ARP (Address Resolution Protocol)* cache is a table maintained locally in each station on a network. ARP cache entries are learned by examining the source information in the ARP packet payload fields, regardless of whether it is an ARP request or response. Thus, when an ARP request is broadcast to all stations on a LAN segment or *VLAN (Virtual LAN)*, every recipient has the opportunity to store the sender's IP and MAC address in their respective ARP cache. The ARP response, being unicast, is normally seen only by the requestor, who stores the sender information in its ARP cache. Newer information always replaces existing content in the ARP cache.

The ARP cache can support 1024 entries, although this size is user-configurable to any value less than 1024. When multiple network interfaces are supported by a device, as is typical of a router, either a single ARP cache is used for all interfaces, or a separate cache is maintained per interface. While the latter approach is useful when network addressing is not unique per interface, this is not the case for Ethernet MAC address assignment so a single ARP cache is employed.

To display the system ARP cache, click **System > Status > ARP Cache** page in the navigation menu.

Table 5: ARP Cache Fields

Field	Description
MAC Address	Displays the physical (MAC) address of the system in the ARP cache.
IP Address	Displays the IP address associated with the system's MAC address.
Interface	Displays the unit, slot, and port number being used for the connection. For non-stacking systems, only the slot and port number is displayed. For units that have a service port, the service port will be listed as Management in this field.

Click **Refresh** to reload the page and refresh the ARP cache view. Click **Clear Entries** to clear all entries from the table. The table will be repopulated as new addresses are learned.

Viewing Inventory Information

Use the **Inventory Information** page to display the switch's Vital Product Data, which is stored in non-volatile memory at the factory.

To display this page, click **System > Summary > Inventory** in the navigation menu.

Table 6: Inventory Information Fields

Field	Description
Management Unit Number	Unit number that corresponds to the stack manager. This field is available only on switches that support stacking.
System Description	The product name of this switch.
Machine Type	The machine type of this switch.
Machine Model	The model within the machine type.
Serial Number	The unique serial number for this switch.
FRU Number	The field replaceable unit number.
Part Number	The manufacturing part number.
Maintenance Level	The identification of the hardware change level.
Manufacturer	The two-octet code that identifies the manufacturer.
Burned In MAC Address	The burned-in universally administered MAC address of this switch.
Software Version	The release version.maintenance number of the code currently running on the switch. For example, if the release is 1, the version is 2 and the maintenance number is 4, the format is "1.2.4."
Operating System	The operating system currently running on the switch.
Network Processing Device	Identifies the network processor hardware.
Additional Packages	A list of the optional software packages installed on the switch, if any.

Click **Refresh** to update the information on the screen.

Viewing MAC Addresses

The **MAC Address Table** keeps track of the MAC addresses that are associated with each port. This table allows the device to forward unicast traffic through the appropriate port. The MAC address table is sometimes called the bridge table or the FDB (forwarding database).

Use this page to display information about entries in the MAC address table. The transparent bridging function uses these entries to determine how to forward a received frame.

To display this page, click **System > Summary > MAC Address Table** in the navigation menu.

To remove dynamically learned FDB table entries, use the **Clear** button. You can clear all of the learned entries, the entries for a specific interface or VLAN, or the range of entries that match the MAC address and mask combination.

Table 7: MAC Address Table Fields

Field	Description
VLAN ID	The VLAN with which the MAC address is associated. A MAC address can be associated with multiple VLANs.
MAC Address	A unicast MAC address for which the switch has forwarding and/or filtering information. The format is a six-byte MAC address, with each byte separated by colons.
Interface	The port where this address was learned. The port identified in this field is the port through which the MAC address can be reached.
Interface Index	The Interface Index of the MIB interface table entry associated with the source port. This value helps identify an interface when using SNMP to manage the device.
Status	Information about the entry and why it is in the table, which can be one of the following: <ul style="list-style-type: none"> • Static – The address has been manually configured and does not age out. • Learned – The address has been automatically learned by the device and can age out when it is not in use. Dynamic addresses are learned by examining information in incoming Ethernet frames. • Management – The burned-in MAC address of the device. • Self – The MAC address belongs to one of the device's physical interfaces. • GMRP Learned – The address was added dynamically by the <u>GARP (Generic Attribute Registration Protocol)</u> Multicast Registration Protocol. • Other – The address was added dynamically through an unidentified protocol or method. • Unknown – The device cannot determine the status of the entry.

Click **Refresh** to update the information on the screen.

Viewing System Resources

Use the **System Resources** page to display the following memory information for the switch:

- Free memory
- Allocated memory
- CPU utilization by task
- Total CPU utilization at the following intervals:
 - Five seconds

- One minute
- Five minutes

To display this page, click **System > Status > Resource Status** in the navigation menu.

Table 8: System Resource Status Fields

Field	Description
Free Memory	Displays the available Free Memory on the switch.
Alloc Memory	Displays the allocated Memory for the switch.
Task Id	Displays the Id of running tasks.
Task Name	Displays the name of the running tasks.
CPU Utilization Report	Displays the Total CPU Utilization in terms of percentage. Total CPU Utilization is shown in the following intervals: <ul style="list-style-type: none"> • 5 seconds • 60 seconds • 300 seconds

To display the **Resource Configuration** page, click **System > Status > Resource Configuration** in the navigation menu.

Table 9: System Resource Configuration Fields

Field	Description
Rising Threshold	The CPU Rising utilization threshold in percentage. A zero percent threshold indicates CPU Utilization Notification feature is disabled.
Rising Threshold Interval	The CPU Rising threshold interval in seconds. The time interval is configured in multiples of five. A time interval of zero seconds indicates CPU Utilization Notification feature is disabled.
Falling Threshold	The CPU Falling utilization threshold in percentage. Configuration of this field is optional. If configured, the Falling threshold value must be equal to or less than the Rising threshold value. If not configured, it takes the same value as the Rising threshold.
Falling Threshold Interval	The CPU Falling threshold interval in seconds. Configuration of this field is optional. If configured, the Falling interval value must be equal to or less than the Rising interval value. If not configured, it takes the same value as the Rising interval. The time interval is configured in multiples of five.
Free Memory Threshold	The CPU Free Memory threshold in kilobytes. A zero threshold value indicates CPU Free Memory Notification feature is disabled.

Click **Submit** to send the updated configuration to the switch. Click **Refresh** to update the page with the most current information. Click **Cancel** to exit the page.

Defining General Device Information

The **Configuration** submenu in the **System** menu contains links to pages that allow you to configure device parameters. The **Configuration** folder contains links to the following features:

- [System Description](#) on page 39
- [Switch Configuration](#) on page 40
- [IP Address Conflict Detection](#) on page 41
- [Service Port IPv4](#) on page 42
- [Service Port IPv6 Neighbors](#) on page 44
- [Network Port DHCPv6 Client Statistics](#) on page 45
- [Network Connectivity Configuration](#) on page 46
- [HTTP Configuration](#) on page 52
- [Network Port IPv6 Neighbors](#) on page 47
- [DHCP Client Options](#) on page 49
- [Telnet Session](#) on page 50
- [Serial Port Configuration](#) on page 51
- [User Accounts](#) on page 57
- [User Domain Name](#) on page 60
- [Authentication Server Users](#) on page 59
- [Authentication List Summary](#) on page 65
- [Select Authentication List](#) on page 66
- [Line Password](#) on page 69
- [Enable Password](#) on page 70
- [Password Rules](#) on page 70
- [Denial of Service](#) on page 72

System Description

After a successful login, the **System Description** window opens. Use this page to configure and view general device information.

To display this page, click **System > Summary > Description** in the navigation menu.

Table 10: System Description Fields

Field	Description
System Description	The product name of this switch.
System Name	Enter the name you want to use to identify this switch. You may use up to 31 alphanumeric characters. The factory default is blank.
System Location	Enter the location of this switch. You may use up to 31 alphanumeric characters. The factory default is blank.
System Contact	Enter the contact person for this switch. You may use up to 31 alphanumeric characters. The factory default is blank.
IP Address	The IP address assigned to the network interface. To change the IP address, see Network Connectivity Configuration on page 46.
Service Port IP Address	The IP address assigned to the service port. The service port provides remote management access to the device. Traffic on this port is segregated from operational network traffic on the switch ports and cannot be switched or routed to the operational network.

Table 10: System Description Fields (continued)

Field	Description
System Object ID	The base object ID for the switch's enterprise MIB.
System Up Time	Displays the number of days, hours, and minutes since the last system restart.
Current SNTP Synchronized Time	Displays currently synchronized <i>SNTP (Simple Network Time Protocol)</i> time in UTC. If no SNTP server has been configured and the time is not synchronized, this field displays <i>Not Synchronized</i> . To specify an SNTP server, see Configuring SNTP Settings on page 97.
MIBs Supported	Displays the list of MIBs supported by the management agent running on this switch.

Defining System Information

The system parameters are applied, and the device is updated.



Note

If you want the switch to retain the new values across a power cycle, you must perform a save.

- 1 Open the **System Description** page.
- 2 Define the following fields: **System Name**, **System Contact**, and **System Location**.
- 3 Scroll to the bottom of the page and click **Submit**.

Switch Configuration

IEEE 802.3x flow control works by pausing a port when the port becomes oversubscribed and dropping all traffic for small bursts of time during the congestion condition. This can lead to high-priority and/or network control traffic loss. When 802.3x flow control is enabled, lower speed switches can communicate with higher speed switches by requesting that the higher speed switch refrains from sending packets. Transmissions are temporarily halted to prevent buffer overflows.

To display the **Switch Configuration** page, click **System > Basic Configuration > Switch** in the navigation menu.

Table 11: Switch Configuration Fields

Field	Description
IEEE 802.3x Flow Control Mode	<p>The 802.3x flow control mode on the switch. IEEE 802.3x flow control works by pausing a port when the port becomes oversubscribed. This allows lower-speed switches to communicate with higher-speed switches. A lower-speed or congested switch can send a PAUSE frame requesting that the peer device refrain from sending packets. Transmissions are temporarily halted to prevent buffer overflows. The options are as follows:</p> <ul style="list-style-type: none"> • Disabled: The switch does not send PAUSE frames if the port buffers become full. • Enabled: The switch can send PAUSE frames to a peer device if the port buffers become full.
MAC Address Aging Interval	<p>The MAC address table (forwarding database) contains static entries, which never age out, and dynamically-learned entries, which are removed if they are not updated within a given time. Specify the number of seconds a dynamic address should remain in the MAC address table after it has been learned.</p>

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

IP Address Conflict Detection

Use the **IP Address Conflict Detection** page to run the IP Address Conflict Detection tool, which detects IP address conflicts for IPv4 addresses. When a conflict is detected, the switch updates the status on the page, generates an [*SNMP \(Simple Network Management Protocol\)*](#) trap, and a logs a message noting the conflict.

To display this page, click **System > Utilities > IP Address Conflict** in the navigation menu.

Table 12: IP Address Conflict Detection Fields

Field	Description
IP Address Conflict Currently Exists	Shows whether an address conflict has been detected since status was last reset.
Last Conflicting IP Address	The IP address of the interface that was last found to be conflicting. This field displays only if a conflict has been detected since the switch was last rebooted.
Last Conflicting MAC Address	The MAC address of the remote host associated with the IP address that was last found to be conflicting. This field displays only if a conflict has been detected since the switch was last rebooted.
Time Since Conflict Detected	The time elapsed – displayed in days, hours, minutes, seconds – since the last conflict was detected (provided a reboot did not occur in the meantime). This field displays only if a conflict has been detected since the switch was last rebooted.

To run the tool and check for possible address conflicts, click **Run Conflict Detection**. If the conflict detection status is true, click **Reset Conflict Detection Status** to clear the information and run the tool again.

Service Port IPv4

Some platforms have a built-in service port that can serve as a dedicated network management interface. For systems that have the service port, the **Service Port IPv4 Configuration** page allows you to configure network information for the switch.

To access this page, click **System > Connectivity > Service Port IPv4** in the navigation menu.

Table 13: Service Port IPv4 Configuration Fields

Field	Description
IPv4 Fields: These display IPv4 configuration information.	
Service Port Configuration Protocol	Specify what the switch should do following power-up. The factory default is None. The options are as follows: <ul style="list-style-type: none"> • None: Do not send any requests following power-up. • BootP: Transmit a Bootp request. • DHCP: Transmit a <i>DHCP (Dynamic Host Configuration Protocol)</i> request.
IP Address	The IP address of the network interface. The factory default value is 0.0.0.0. Each part of the IP address must start with a number other than zero. For example, IP addresses 001.100.192.6 and 192.001.10.3 are not valid.
Subnet Mask	The IP subnet mask for the interface. The factory default value is 0.0.0.0.
Default Gateway	The default gateway for the IP interface. The factory default value is 0.0.0.0.
Burned-in MAC Address	This read-only field displays the MAC address that is burned-in to the network card at the factory. This MAC address is used for in-band connectivity if you choose not to configure a locally administered address.
IPv6 Fields: If the system supports IPv6, these fields display IPv6 configuration information.	
IPv6 Mode	Enables or disables IPv6 mode on the interface.
Service Port Configuration Protocol	Specify what the switch should do following power-up. The factory default is None. The options are as follows: <ul style="list-style-type: none"> • DHCP: Transmit a DHCP request. • None: Do not send any requests following power-up.
IPv6 Stateless Address AutoConfig Mode	Enables or disables the IPv6 stateless address autoconfiguration on the management port. The factory default is None.
Change IPv6 Gateway	Select the checkbox to configure an IPv6 address.
IPv6 Gateway	Enter the IPv6 gateway address (do not include a prefix).
Add/Delete IPv6 Address	Select to add or remove IPv6 Addresses. The fields New IPv6 Address and EUI Flag are visible when Add is selected from this menu.
New IPv6 Address	Displays when Add IPv6 Address is selected. Adds IPv6 address.
EUI Flag	Displays when Add IPv6 Address is selected. Sets the EUI flag while configuring a new IPv6 address when TRUE is selected. The default is FALSE .
IPv6 Addresses	Displays IPv6 addresses.
Default Routers	Displays the address(es) entered in the IPv6 Gateway field.

To renew the IPv4 address learned from a DHCP server on the service port, click **Renew DHCP IPv4 Address**.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Service Port IPv6

Use the **Service Port Configuration** page to configure IPv6 network information on the service port. The service port is a dedicated Ethernet port for out-of-band management of the device. Traffic on this port is segregated from operational network traffic on the switch ports and cannot be switched or routed to the operational network.

To access this page, click **System > Connectivity > Service Port IPv6** in the navigation menu.

Table 14: Service Port IPv6 Configuration Fields

Field	Description
IPv6 Mode	Enables or disables IPv6 mode on the interface.
Service Port Configuration Protocol	Specify whether the device should attempt to acquire network information from a DHCPv6 server. Selecting None disables the DHCPv6 client on the service port.
IPv6 Stateless Address AutoConfig Mode	Sets the IPv6 stateless address autoconfiguration mode on the service port. <ul style="list-style-type: none"> • Enabled: The service port can acquire an IPv6 address through IPv6 Neighbor Discovery Protocol (NDP) and through the use of Router Advertisement messages. • Disabled: The service port will not use the native IPv6 address autoconfiguration features to acquire an IPv6 address.
DHCPv6 Client DUID	The client identifier used by the DHCPv6 client (if enabled) when sending messages to the DHCPv6 server.
Static IPv6 Addresses	Lists the manually configured static IPv6 addresses on the service port interface. Use the buttons available in this table to perform the following tasks: <ul style="list-style-type: none"> • To add an entry to the list, click the + (plus) button to open the Add IPv6 Address dialog and provide the following: <ul style="list-style-type: none"> • New IPv6 Address: Specify the IPv6 address to add to the service port interface. • EUI Flag: Select this option to enable the Extended Universal Identifier (EUI) flag for IPv6 address, or clear the option to omit the flag. • To delete an entry from the list, click the - (minus) button associated with the entry to remove. • To delete all entries from the list, click the - (minus) button in the heading row.
Dynamic IPv6 Addresses	Lists the IPv6 addresses on the service port interface that have been dynamically configured through IPv6 autoconfiguration or DHCPv6.
Default IPv6 Routers	Lists the IPv6 address of each default router that has been automatically configured through IPv6 router discovery.

To renew the IPv6 address learned from a DHCP server on the service port, click **Renew DHCP IPv6 Address**.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Service Port IPv6 Neighbors

The **Service Port IPv6 Neighbors** page provides information about IPv6 neighbors the device has discovered through the service port by using the Neighbor Discovery Protocol (NDP). The manually configured static service port IPv6 neighbors are also displayed.

To display this page, click **System > Connectivity > Service Port IPv6 Neighbors**

Table 15: Service Port IPv6 Neighbors Fields

Field	Description
IPv6 Addresses	Displays the IP address of the neighbor.
MAC Address	Displays the MAC address of the neighbor.
Type	The type of the neighbor entry, which is one of the following: <ul style="list-style-type: none"> • Static: The neighbor entry is manually configured. • Dynamic: The neighbor entry is dynamically resolved. • Local: The neighbor entry is local. • Other: The neighbor entry is unknown.
Is Router	Whether the neighbor is a router. If the neighbor is a router, the value is TRUE. If the neighbor is not a router, the value is FALSE.
Neighbor State	Specifies the state of the neighbor cache entry. Following are the states for dynamic entries in the IPv6 neighbor discovery cache: <ul style="list-style-type: none"> • Reachable: Positive confirmation was received within the last Reachable Time milliseconds that the forward path to the neighbor was functioning properly. While in REACH state, the device takes no special action as packets are sent. • Stale: More than ReachableTime milliseconds have elapsed since the last positive confirmation was received that the forward path was functioning properly. While in STALE state, the device takes no action until a packet is sent. • Delay: More than ReachableTime milliseconds have elapsed since the last positive confirmation was received that the forward path was functioning properly. A packet was sent within the last DELAY_FIRST_PROBE_TIME seconds. If no reachability confirmation is received within DELAY_FIRST_PROBE_TIME seconds of entering the DELAY state, send a neighbor solicitation message and change the state to PROBE. • Probe: A reachability confirmation is actively sought by resending neighbor solicitation messages every RetransTimer milliseconds until a reachability confirmation is received. • Unknown: The reachability status cannot be determined.
Last Updated	Displays the time since the address was confirmed to be reachable.

Use the buttons to perform the following tasks:

- To add service port static IPv6 neighbor entry, click **Add** and configure the desired settings.
- To remove service port static IPv6 neighbor entries, select each static neighbor entry to remove and click **Remove**.

After you click **Add** or **Edit**, a window opens and allows you to configure Service Port IPv6 Neighbor settings.

You can configure the IPv6 address and the MAC address.

Table 16: Add Service Port IPv6 Neighbor Fields

Field	Description
IPv6 Address	IP address of the neighbor.
MAC Address	MAC address of the neighbor.

Network Port DHCPv6 Client Statistics

The **Network Port DHCPv6 Client Statistics** page displays the DHCPv6 client statistics values for the network interface. The DHCPv6 client on the device exchanges several different types of UDP messages with one or more network DHCPv6 servers during the process of acquiring address, prefix, or other relevant network configuration information from the server. The values indicate the various counts that have accumulated since they were last cleared.

To display this page, click **System > Statistics > Statistics > Network DHCPv6**.

Table 17: Network Port DHCPv6 Client Statistics Fields

Field	Description
Advertisement Packets Received	Number of DHCPv6 advertisement messages received from one or more DHCPv6 servers in response to the client's solicit message.
Reply Packets Received	Number of DHCPv6 reply messages received from one or more DHCPv6 servers in response to the client's request message.
Received Advertisement Packets Discarded	Number of DHCPv6 advertisement messages received from one or more DHCPv6 servers to which the client did not respond.
Received Reply Packets Discarded	Number of DHCPv6 reply messages received from one or more DHCPv6 servers to which the client did not respond.
Malformed Packets Received	Number of messages received from one or more DHCPv6 servers that were improperly formatted.
Total Packets Received	Total number of messages received from all DHCPv6 servers.
Solicit Packets Transmitted	Number of DHCPv6 solicit messages the client sent to begin the process of acquiring network information from a DHCPv6 server.
Request Packets Transmitted	Number of DHCPv6 request messages the client sent in response to a DHCPv6 server's advertisement message.
Renew Packets Transmitted	Number of renew messages the DHCPv6 client has sent to the server to request an extension of the lifetime of the information provided by the server. This message is sent to the DHCPv6 server that originally assigned the addresses and configuration information.
Rebind Packets Transmitted	Number of rebind messages the DHCPv6 client has sent to any available DHCPv6 server to request an extension of its addresses and an update to any other relevant information. This message is sent only if the client does not receive a response to the renew message.

Table 17: Network Port DHCPv6 Client Statistics Fields (continued)

Field	Description
Release Packets Transmitted	Number of release messages the DHCPv6 client has sent to the server to indicate that it no longer needs one or more of the assigned addresses.
Total Packets Transmitted	Total number of messages sent to all DHCPv6 servers.

Click **Clear Counters** to clear all the statistics counters, resetting all switch summary and detailed statistics to default values.

Network Connectivity Configuration

The network interface is the logical interface used for in-band connectivity with the switch via any of the switch's front panel ports. The configuration parameters associated with the switch's network interface do not affect the configuration of the front panel ports through which traffic is switched or routed.

The **Network Connectivity** page allows you to change the IPv4 and IPv6 information using the web interface. To access this page, click **System > Connectivity > IPv4** or **IPv6** in the navigation menu.

Table 18: Network Connectivity Configuration for IPv4 and IPv6 Fields

Field	Description
IPv4 Fields	
Network Configuration Protocol	Specify what the switch should do following power-up. The factory default is None. The options are as follows: <ul style="list-style-type: none"> • None: Do not send any requests following power-up. • Bootp: Transmit a BOOTP request. • DHCP: Transmit a DHCP request.
DHCP Client Identifier	The DHCP Client Identifier (Option 61) is used by DHCP clients to specify their unique identifier. DHCP servers use this value to index their database of address bindings. This value is expected to be unique for all clients in an administrative domain. The Client Identifier string will be displayed beside the checkbox once DHCP is enabled on the port on which the Client Identifier option is selected. This web page will need to be refreshed once this change is made.
IP Address	The IP address of the network interface. The factory default value is 0.0.0.0. Each part of the IP address must start with a number other than zero. For example, IP addresses 001.100.192.6 and 192.001.10.3 are not valid.
Subnet Mask	The IP subnet mask for the interface. The factory default value is 0.0.0.0.
Default Gateway	The default gateway for the IP interface. The factory default value is 0.0.0.0.
MAC Address Type	Specify whether the burned-in or the locally administered MAC address should be used for in-band connectivity. The factory default is to use the burned-in MAC address
Burned-in MAC Address	This read-only field displays the MAC address that is burned-in to the network card at the factory. This MAC address is used for in-band connectivity if you choose not to configure a locally administered address.

Table 18: Network Connectivity Configuration for IPv4 and IPv6 Fields (continued)

Field	Description
Locally Administered MAC Address	Specifies a locally administered MAC address for in-band connectivity instead of using the burned-in universally administered MAC address. In addition to entering an address in this field, you must also set the MAC address type to locally administered. Enter the address as twelve hexadecimal digits (6 bytes) with a colon between each byte. Bit 1 of byte 0 must be set to a 1 and bit 0 to a 0. In other words, byte 0 must have a value between x'40' and x'7F'.
Management VLAN ID	Specify the management VLAN ID of the switch. It may be configured to any value in the range of 1 to 4093. The management VLAN is used for management of the switch. This field is configurable for administrative users and read-only for other users.
IPv6 Fields	
IPv6 Mode	Enable or disable IPv6 mode.
Network Configuration Protocol	Enable or Disable DHCPv6 Client protocol on the management port. The factory default is None .
IPv6 Stateless Address AutoConfig Mode	Enables or Disables the IPv6 stateless address autoconfiguration on the management port. The factory default is None .
DHCPv6 Client DUID	Displays the Client Identifier used by DHCPv6 Client when sending messages to the DHCPv6 Server. This entry displays only if IPv6 Network Configuration Protocol is set to DHCP.
Change IPv6 Gateway	Select the checkbox to configure an IPv6 Address.
IPv6 Gateway	Enter the IPv6 gateway address (do not include a prefix).
Add/Delete IPv6 Address	Select to add or remove IPv6 Addresses. The fields New IPv6 Address and EUI Flag are visible when Add is selected.
New IPv6 Address	Displays when Add IPv6 Address is selected. Adds IPv6 address.
EUI Flag	Displays when Add IPv6 Address is selected. Sets the EUI flag while configuring a new IPv6 address when selected. The Default is deselected.
Dynamic IPv6 Addresses	Lists the IPv6 addresses on the network interface that have been dynamically configured through IPv6 autoconfiguration or DHCPv6.
IPv6 Addresses	Displays the configured IPv6 addresses.
Default IPv6 Routers	Displays the default IPv6 Router address(es).

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Click **Renew DHCP IPv4 Address** to force the interface to release the current DHCP-assigned information and submit a request for new information.

Network Port IPv6 Neighbors

The **Network Port IPv6 Neighbors** page provides information about IPv6 neighbors the device has discovered through the network interface by using the Neighbor Discovery Protocol (NDP) and the manually configured static network port IPv6 neighbors.

To access this page, click **System > Connectivity > IPv6 Neighbors**.

Table 19: Network Port IPv6 Neighbors Fields

Field	Description
IPv6 Address	Displays the IP address of the neighbor.
MAC Address	Displays the MAC address of the neighbor.
IS Router	Whether the neighbor is a router. If the neighbor is a router, the value is TRUE. If the neighbor is not a router, the value is FALSE.
Type	The type of the neighbor entry, which is one of the following: <ul style="list-style-type: none"> • Static: The neighbor entry is manually configured. • Dynamic: The neighbor entry is dynamically resolved. • Local: The neighbor entry is a local entry. • Other: The neighbor entry is an unknown entry.
Neighbor State	Specifies the state of the neighbor cache entry. Following are the states for dynamic entries in the IPv6 neighbor discovery cache: <ul style="list-style-type: none"> • Reachable: Positive confirmation was received within the last Reachable Time milliseconds that the forward path to the neighbor was functioning properly. While in REACH state, the device takes no special action as packets are sent. • Stale: More than ReachableTime milliseconds have elapsed since the last positive confirmation was received that the forward path was functioning properly. While in STALE state, the device takes no action until a packet is sent. • Delay: More than ReachableTime milliseconds have elapsed since the last positive confirmation was received that the forward path was functioning properly. A packet was sent within the last DELAY_FIRST_PROBE_TIME seconds. If no reachability confirmation is received within DELAY_FIRST_PROBE_TIME seconds of entering the DELAY state, send a neighbor solicitation message and change the state to PROBE. • Probe: A reachability confirmation is actively sought by resending neighbor solicitation messages every RetransTimer milliseconds until a reachability confirmation is received. • Unknown: The reachability status cannot be determined.
Last Updated	Displays the time since the address was confirmed to be reachable.

Use the buttons to perform the following tasks:

- To add network port static IPv6 neighbor entry, click **Add** and configure the desired settings.
- To remove network port static IPv6 neighbor entries, select each static neighbor entry to remove and click **Remove**.

After you click **Add** or **Edit**, a window opens and allows you to configure Network Port IPv6 Neighbor settings.

You can configure the IPv6 Address and the MAC Address.

Table 20: Add Network Port IPv6 Neighbor Fields

Field	Description
IPv6 Address	Use this field to enter the IP address of the neighbor.
MAC Address	Use this field to enter the MAC address of the neighbor.

DHCP Client Options

Use the **DHCP Client Options** page to configure DHCP client settings on the system.

To access this page, click **System > Connectivity > DHCP Client Options** in the navigation menu.

Table 21: DHCP Client Options Fields

Field	Description
DHCP Vendor Class ID Mode	Enables/Disables the vendor class identifier mode.
DHCP Vendor Class ID String	The string added to DHCP requests as Option-60: the Vendor Class Identifier option.

System Connectivity

Use the **System Connectivity** page to control access to the management interface by administratively enabling or disabling various access methods.

To display this page, click **System > Management Access > System** in the navigation menu.

Table 22: System Connectivity Configuration Fields

Field	Description
HTTP Admin Mode	Enables or disables the HTTP administrative mode. When this mode is enabled, the device management interface can be accessed through a web browser using the HTTP protocol.
Telnet Server Admin Mode	Enables or disables the Telnet administrative mode. When this mode is enabled, the device CLI can be accessed through the telnet port. Disabling this mode disconnects all existing telnet connections and shuts down the telnet port in the device.
Telnet—Allow New Sessions	Enables or disables new Telnet sessions. When this option is disabled, the system does not accept any new telnet sessions, but existing telnet sessions are unaffected.
Outbound Telnet—Allow New Sessions	Enables or disables new outbound telnet sessions. When this option is disabled, initiating telnet sessions from the system is not allowed.
HTTPS Admin Mode	Enables or disables the administrative mode of secure HTTP. When this mode is enabled, the device management interface can be accessed through a web browser using the HTTPS protocol.
SSH Admin Mode	Enables or disables the administrative mode of SSH. When this mode is disabled, all existing SSH connections remain connected until timed-out or logged out, but new SSH connections cannot be established.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Telnet Session

Telnet is a terminal emulation TCP/IP protocol. ASCII terminals can be virtually connected to the local device through a TCP/IP protocol network. Telnet is an alternative to a local login terminal where a remote login is required.

The switch supports up to five simultaneous Telnet sessions. All CLI commands can be used over a Telnet session.

The **Telnet Session Configuration** page allows you to control inbound Telnet settings on the switch. Inbound Telnet sessions originate on a remote system and allow a user on that system to connect to the switch CLI.

To display this page, click **System > Management Access > Telnet** in the navigation menu.

Table 23: Telnet Session Configuration Fields

Field	Description
Admin Mode	Enables or disables the Telnet administrative mode. When enabled, the device may be accessed through the Telnet port (23). Disabling this mode value disconnects all existing Telnet connections and shuts down the Telnet port in the device.
Telnet Port	The TCP port number on which the Telnet server listens for requests. Existing Telnet login sessions are not affected by a change in this value, although establishment of any new Telnet sessions must use the new port number. Before changing this value, verify that the desired port number is not currently being used by any other service.
Session Timeout (Minutes)	Specify how many minutes of inactivity should occur on a Telnet session before the session is logged off. You may enter any number from 1 to 160. The factory default is 5. When you change the timeout value, the new value is applied to all active and inactive sessions immediately. Any sessions that have been idle longer than the new timeout value are disconnected immediately.
Maximum Number of Sessions	From the drop-down menu, select how many simultaneous Telnet sessions to allow. The maximum is 4, which is also the factory default. A value of 0 indicates that no outbound Telnet session can be established.
Allow New Sessions	Controls whether to allow new Telnet sessions: <ul style="list-style-type: none"> • Yes: Permits new Telnet sessions until the maximum number allowed is reached. • No: New Telnet sessions will not be allowed, but existing sessions are not disconnected.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Outbound Telnet Configuration

The **Outbound Telnet Configuration** page displays the current value of the outbound Telnet settings on the device. An outbound Telnet session is a Telnet session initiated from the CLI of the device to the Telnet client on a remote device.

To display this page, click **System > Management Access > Outbound Telnet** in the navigation menu.

Table 24: Outbound Telnet Configuration Fields

Field	Description
Allow New Sessions	Controls whether new outbound Telnet sessions are allowed. Setting this value to Disable disallows any new outbound Telnet sessions from starting (although existing Telnet sessions are unaffected).
Maximum Number of Sessions	The maximum number of allowed outbound Telnet sessions from the device simultaneously.
Session Timeout	Outbound Telnet session inactivity timeout value, in minutes. An outbound Telnet session is closed automatically if there is no activity within the configured amount of time.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Serial Port Configuration

The **Serial Port Configuration** page allows you to change the switch's serial port settings. In order for a terminal or terminal emulator to communicate with the switch, the serial port settings on both devices must be the same. Some settings on the switch cannot be changed.

To view or configure the serial port settings on the switch, click **System > Management Access > Serial** in the navigation menu.

Table 25: Serial Port Fields

Field	Description
Serial Port Time Out (Minutes)	Indicates how many minutes of inactivity should occur on a serial port connection before the switch closes the connection. Enter a number between 0 and 160. The factory default is 5. Entering 0 disables the timeout.
Baud Rate (bps)	Select the default baud rate for the serial port connection from the menu. The factory default is 9600 baud.
Character Size (Bits)	The number of bits in a character. This value is always 8.
Parity	The parity method used on the serial port. This value is always None.
Stop Bits	The number of stop bits per character. This value is always 1.
Flow Control	Whether hardware flow control is enabled or disabled. It is always disabled.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

CLI Banner Configuration

Use the **CLI Banner Configuration** page to configure a message that appears before the user prompt as a Pre-login banner. The message configured shows up on Telnet, SSH and Console connections.

To access this page, click **System > Management Access > CLI Banner** in the navigation menu.

Table 26: CLI Banner Configuration Fields

Field	Description
CLI Banner Message	Text area for creating, viewing, or updating the CLI banner message. To create the CLI banner message, type the desired message in the text area. If you reach the end of the line, the text wraps to the next line. The line might not wrap at the same location in the CLI. To create a line break in the message, press [Enter] . The line break in the text area will be at the same location in the banner message when viewed through the CLI.
Clear (Button)	Clears the CLI banner message from the device. After clicking Clear , you must confirm the action. You can also clear the CLI banner by deleting the text in the CLI Banner Message field and clicking Submit .

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

HTTP Configuration

Use the **HTTP Configuration** page to configure the HTTP server settings on the system.

To access this page, click **System > Management Access > HTTP** in the navigation menu.

Table 27: HTTP Configuration Fields

Field	Description
HTTP Admin Mode	Enables or disables the Administrative Mode of HTTP. The currently configured value is shown when the web page is displayed. The default value is Enable. If you disable the HTTP admin mode, access to the web interface is limited to secure HTTP, which is disabled by default.
Java Mode	Enables or disables the web Java Mode. This applies to both secure and un-secure HTTP connections. The currently configured value is shown when the web page is displayed. The default value is Enable.
HTTP Port	The TCP port number on which the HTTP server listens for requests. Existing HTTP login sessions are closed whenever this value is changed. All new HTTP sessions must use the new port number. Note: Before changing this value, verify that the desired port number is not currently being used by any other service.
HTTP Session Soft Timeout	Sets the inactivity timeout for HTTP sessions. The value must be in the range of 1 to 60 minutes. A value of zero corresponds to an infinite timeout. The default value is 5 minutes. The currently configured value is shown when the web page is displayed.

Table 27: HTTP Configuration Fields (continued)

Field	Description
HTTP Session Hard Timeout	This field is used to set the hard timeout for HTTP sessions. This timeout is unaffected by the activity level of the session. The value must be in the range of 1 to 168 hours. A value of 0 corresponds to an infinite timeout. The default value is 24 hours. The currently configured value is shown when the web page is displayed.
Maximum Number of HTTP Sessions	Sets the maximum allowable number of HTTP sessions. The value must be in the range of 0 to 16. The default value is 16. The currently configured value is shown when the web page is displayed.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

HTTPS Configuration

Use the **HTTPS Configuration** page to view and modify the Secure HTTP settings on the device. HTTPS increases the security of web-based management by encrypting communication between the administrative system and the device.

To access this page, click **System > Management Access > HTTPS** in the navigation menu.

Table 28: HTTPS Configuration Fields

Field	Description
HTTPS Admin Mode	Enables or disables the HTTPS administrative mode. When this mode is enabled, the device can be accessed through a web browser using the HTTPS protocol.
TLS Version 1	Enables or disables <u>TLS</u> version 1.0. When this option is enabled, communication between the web browser on the administrative system and the web server on the device is sent through TLS 1.0.
SSL Version 3	Enables or disables <u>SSL (Secure Socket Layer)</u> version 3.0. When this option is enabled, communication between the web browser on the administrative system and the web server on the device is sent through SSL 3.0. SSL must be administratively disabled while downloading an SSL certificate file from a remote server to the device.
HTTPS Port	The TCP port number that HTTPS uses.
HTTPS Session Soft Time Out (Minutes)	HTTPS session inactivity timeout value. A logged-in user that does not exhibit any HTTPS activity for this amount of time is automatically logged out of the HTTPS session.
HTTPS Session Hard Time Out (Hours)	HTTPS session hard timeout value. A user connected to the device via an HTTPS session is automatically logged out after this amount of time regardless of the amount of HTTPS activity that occurs.
Maximum Number of HTTPS Sessions	The maximum number of HTTPS sessions that can be connected to the device simultaneously.

Table 28: HTTPS Configuration Fields (continued)

Field	Description
Certificate Status	The status of the SSL certificate generation process. <ul style="list-style-type: none"> • Present: The certificate has been generated and is present on the device • Absent: Certificate is not available on the device • Generation in Progress: An SSL certificate is currently being generated.
 Download Certificates (Button)	Allows you to download an SSL certificate file from a remote system to the device. Note that to download SSL certificate files, SSL must be administratively disabled.
 Generate Certificate (Button)	Generates an SSL certificate to use for secure communication between the web browser and the embedded web server on the device.
 Delete Certificates (Button)	Deletes the SSL certificate. This button is available only if an SSL certificate is present on the device.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

SSH Configuration

Use the **SSH Configuration** page to view and modify the [SSH \(Secure Shell\)](#) server settings on the device. SSH is a network protocol that enables access to the CLI management interface by using an SSH client on a remote administrative system. SSH is a more secure access method than [Telnet](#) because it encrypts communication between the administrative system and the device. This page also allows you to download or generate SSH host keys for secure CLI-based management.

To access this page, click **System > Management Access > SSH** in the navigation menu.

Table 29: SSH Configuration Fields

Field	Description
SSH Admin Mode	Enables or disables the SSH server administrative mode. When this mode is enabled, the device can be accessed by using an SSH client on a remote system.
SSH Port	The TCP (Transmission Control Protocol) port number on which the SSH server listens for requests. Existing SSH login sessions are not affected by a change in this value, although establishment of any new SSH sessions must use the new port number. Before changing this value, verify that the desired port number is not currently being used by any other service.
SSH Version 1	When this option is selected, the SSH server on the device can accept connections from an SSH client using SSH-1 protocol. If the option is clear, the device does not allow connections from clients using the SSH-1 protocol.

Table 29: SSH Configuration Fields (continued)

Field	Description
SSH Version 2	When this option is selected, the SSH server on the device can accept connections from an SSH client using SSH-2 protocol. If the option is clear, the device does not allow connections from clients using the SSH-2 protocol.
SSH Connections Currently in Use	The number of active SSH sessions between remote SSH clients and the SSH server on the device.
Maximum number of SSH Sessions Allowed	The maximum number of SSH sessions that may be connected to the device simultaneously.
SSH Session Timeout (minutes)	The SSH session inactivity timeout value. A connected user that does not exhibit any SSH activity for this amount of time is automatically disconnected from the device.
RSA Key Status	The status of the SSH-1 <i>RSA</i> key file or SSH-2 RSA key file (PEM Encoded) on the device, which might be Present, Absent, or Generation in Progress.
DSA Key Status	The status of the SSH-2 <i>DSA</i> key file (PEM Encoded) on the device, which might be Present, Absent, or Generation in Progress.
Download Certificates (Button) 	Use this button to download an SSH-1 RSA, SSH-2 RSA, or SSH-2 DSA key file from a remote system to the device. After you click the button, a Download Certificate window opens. Select the file type to download, browse to the location on the remote system, and select the file to upload. Then, click Begin Transfer . The Status field provides information about the file transfer.
Generate Certificate (Button) 	Use this button to manually generate an RSA key or DSA key on the device.
Delete Certificates (Button) 	Use this button to delete an RSA key or DSA key that has been downloaded to the device or manually generated on the device.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Management Access Control and Administration List

Use the **Management Access List Configuration** page to create and configure a management access list to help secure access to the switch management features. The Management Access Control and Administration List (MACAL) feature is used to ensure that only known and trusted devices are allowed to remotely manage the switch via TCP/IP.

This page provides the capability to add, edit, and remove MACALs. MACALs can be applied only to in-band ports and cannot be applied to the service port.

To access this page, click **System > Management Security > Access Profile** in the navigation menu.

Note

Profile rules cannot be added or modified when a profile is active. To add or edit a profile, the Active Profile field must be set to **None**.



- To add a new MACAL, click **Add**. The **Add Profile Rule** dialog box opens. Specify the rule criteria in the available fields.
- To edit an existing rule, select the appropriate checkbox or click the row to select the account and click **Edit**. The **Edit Profile Rule** box opens. Modify the rule criteria as needed.
- To remove a Profile Rule, select one or more table entries and click **Remove** to delete the selected entries.

Table 30: Management Access List Configuration Fields

Field	Description
Access Profile	Profile name for the Management Access Control list. One user defined Access Profile can be created.
Active Profile	Currently enabled profile name.
Packets Filtered	The number of packets filtered due to matching a rule in the MACAL.
Interface	The port/interface or trunk ID.
Management Method	The types of action will be taken on access control list. <ul style="list-style-type: none"> • Permit: To allow conditions for the management access list. • Deny: To deny conditions for the management access list. In the Add or Edit Profile Rule dialog, this is specified using the Action field.
Source IP Address	IP Address of device which needs to permit or deny in the management access list.
Subnet Mask	Specifies the network mask of the source IP address.
VLAN	The VLAN ID.
Port Channel	Port channels, also known as LAGs (Link Aggregation Groups), allow one or more full-duplex Ethernet links of the same speed to be aggregated together.
Service	The type of service to permit or deny: <ul style="list-style-type: none"> • ANY • Telnet • HTTP • HTTPS • SNMP • SSH • TFTP • SNTP • JAVA
Priority	Priority for the rule. Duplicates are not allowed.

User Accounts

By default, the switch contains two user accounts:

- admin, with 'Read/Write' privileges
- guest, with 'Read Only' privileges

Both of these accounts have blank passwords by default. The names are not case sensitive.

If you log on to the switch with the user account that Read/Write privileges (that is, as admin), you can use the **User Accounts** page to assign passwords and set security parameters for the default accounts. You can also add up to five read-only accounts. You can delete all accounts except for the Read/Write account.



Note

Only a user with Read/Write privileges may alter data on this screen, and only one account can exist with Read/Write privileges.

To access this page, click **System > Users > Accounts** in the navigation menu.

This page provides the capability to add, edit, and remove user accounts.

- To add a user, click **Add**. The Add new user dialog box opens. Specify the new account information in the available fields.
- To edit an existing user, select the appropriate checkbox or click the row to select the account and click **Edit**. The Edit existing user dialog box opens. Modify the account information as needed.
- To remove a user, select one or more table entries and click **Remove** to delete the selected entries.

Table 31: User Accounts Fields

Field	Description
User Name	Enter the name you want to give to the new account. (You can only enter data in this field when you are creating a new account.) User names are up to 32 alphanumeric characters in length and are not case sensitive. Valid characters include all the alphanumeric characters and the dash (-) and underscore (_) characters. User name <code>default</code> is not valid. You can change the Read/Write user name from "admin" to something else, but when you click Submit , you must re-authenticate with the new username.
Password	Enter the optional new or changed password for the account. It will not display as it is typed, only asterisks (*) or dots (.) will display depending on the browser used. Passwords must be greater than eight characters and can be up to 64 characters in length, and are case sensitive.
Confirm	Enter the password again, to confirm that you entered it correctly. This field will not display, but will show asterisks (*) or dots (.)
Access Level	Indicates the access or privilege level for this user. The options are: <ul style="list-style-type: none"> • Read Write: The user can view and modify the configuration. • Read Only: The user can view the configuration but cannot modify any fields. • Suspended: The user exists but is not permitted to log on to the device.

Table 31: User Accounts Fields (continued)

Field	Description
Password Override	Identifies the password override complexity status for this user. <ul style="list-style-type: none"> • Enable: The system does not check the strength of the password. • Disable: When configuring a password, it is checked against the Strength Check rules configured for passwords.
Password Expiration	Indicates the date when this user's current password will expire. This is determined by the date the password was created and the number of days specified in the aging Password Aging setting on the Password Management page.
Contained User Group	The associated user groups for the user.
Operational Permissions	The operational task permissions for the user. In addition to the fields described above, the User Groups table will be populated when you click on each row. To configure this user group, click the Add icon in the header row. To remove the user group, click the Reset icon in the row.
Password Strength	Shows the status of password strength check.
Encrypt password	Select this option to encrypt the password before it is stored on the device.

Adding a User Account

Use the following procedure to add a user account. The system supports one Read/Write user and five Read Only users.

- 1 From the **User** menu, select **Add**.
The screen refreshes.
- 2 Enter a username and password for the new user, then re-enter the password in the **Confirm Password** field.
- 3 Click **Submit** to update the switch with the values on this screen.
If you want the switch to retain the new values across a power cycle, you must perform a save.

Changing User Account Information

You cannot add or delete the Read/Write user, but you can change the username and password. To change the password for an existing account or to overwrite the username on an existing account, use the following procedure.

- 1 From the **User** menu, select the user to change.
The screen refreshes.
- 2 Click **Edit** to change the user settings.
This button is only visible when you have selected a user account with 'Read Only' access. You cannot delete the 'Read/Write' user.
- 3 To alter the username or, delete the existing name in the **Username** field and enter the new username.
To change the password, delete any asterisks (*) in the **Password** and **Confirm Password** fields, and then enter and confirm the new password.
- 4 Click **Submit** to update the switch with the values on this screen.
If you want the switch to retain the new values across a power cycle, you must perform a save.

Removing a User Account

Use the following procedure to remove any of the Read Only user accounts.

- 1 From the **User** menu, select the user to remove.
The screen refreshes.
- 2 Click **Remove** to delete the user.
This button is only visible when you have selected a user account with 'Read Only' access. You cannot remove the 'Read/Write' user.

If you want the switch to retain the new values across a power cycle, you must perform a save.

Authentication Server Users

Use the **Authentication Server Users** page to add and remove users from the local authentication server user database. For some security features, such as IEEE 802.1X port-based authentication, you can configure the device to use the locally stored list of usernames and passwords to provide authentication to users instead of using an external authentication server.



Note

The preconfigured users, admin and guest, are assigned to a pre-configured list named defaultList, which you cannot delete. All newly created users are also assigned to the defaultList until you specifically assign them to a different list.

You can create a text file that contains a list of IAS users to add to the database and then download the file to the switch. The following script is an example of an IAS user text file that contains three users:

```
configure
aaa ias-user username client-1
password my-password1
exit
aaa ias-user username client-2
password aa5c6c251fe374d5e306c62496c3bcf6 encrypted
exit
aaa ias-user username client-3
password 1f3ccb1157
exit
```

After the download completes, client-1, client-2, and client-3 are added to the IAS database. The password for client-2 is encrypted.

When Dot1x authentication is enabled on the ports and the authentication method is LOCAL, port access is allowed only to users in this database that provide the correct name and password.

Use the buttons to perform the following tasks:

- To add a new authentication server user, click **Add**.
- To add a user to the local authentication server database, click **Add** and complete the required information.
- To change the password information for an existing user, select the user to update and click **Edit**.
- To delete a user from the database, select each user to delete and click **Remove**.
- To remove all users from the database, click **Clear All Users**.

When **Add** is selected from **Auth Server Users list**, the **Add New User** window opens.

Table 32: Add New Authentication User Fields

Field	Description
User Name	A unique name used to identify this user account. You configure the user name when you add a new user.
Password Required	Select this option to indicate that the user must enter a password to be authenticated. If this option is clear, the user is required only to enter a valid user name.
Password	Specify the password to associate with the user name (if required).
Confirm	Re-enter the password to confirm the entry.
Encrypted	Select this option to encrypt the password before it is stored on the device.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Logged in Sessions

The **Logged In Sessions** page identifies the users that are logged in to the management interface of the device. The page also provides information about their connections.

To access this page, click **System > Users > Sessions** in the navigation menu.

Table 33: Logged in Sessions Fields

Field	Description
ID	The unique ID of the session.
User Name	The name that identifies the user account.
Connection From	The administrative system that is the source of the connection. For remote connections, this field shows the IP address of the administrative system. For local connections through the console port, this field shows the communication standard for the serial connection.
Idle Time	The amount of time in hours, minutes, and seconds that the logged-on user has been inactive.
Session Time	The amount of time in hours, minutes, and seconds since the user logged onto the system.
Session Type	The type of session, which can be Telnet, Serial, SSH, HTTP, or HTTPS.

Click **Refresh** to update the information on the screen.

User Domain Name

Use the **User Domain Name** page to configure the domain name to send to the authentication server, along with the user name and password, to authenticate a user attempting to access the device management interface. Domain name authentication is supported when user authentication is performed by a RADIUS (Remote Authentication Dial In User Service) server or TACACS+ server.

To access this page, click **System > Users > User Domain Name** in the navigation menu.

Table 34: User Domain Name Fields

Field	Description
User Domain Name Mode	The administrative mode of domain name authentication on the device. When enabled, the domain name is included when the user name and password are sent to the authentication server. The domain name can be input by the user in the User Name field on the login screen in a domain-name\username format, or the domain name can be specified in the Domain Name field .
Domain Name	The domain name to send to the authentication server when the user does not provide one in the User Name field during logon. When only the username is provided, the device sends the username as domain-name\username, where domain-name is the string configured in this field. To configure the domain name, click the Edit icon and specify the desired string. To reset the field to its default value, click the Reset icon and confirm the action.

Task Groups

The **Task Group Configuration** page allows you to add, edit, and remove task groups. Task groups allow users to have different permission levels (read, write, execute, debug) at a per-component level. Task-based authorization uses the concept of components/tasks to define permission for commands for a given user.

Users are assigned to User Groups that are, in turn, associated with Task Groups. Each Task Group is then associated with one or more tasks/components. This feature is supported only for users who are authenticated locally via the web interface.

To access this page, click **System > Users > Task Groups** in the navigation menu.

Table 35: Task Group Configuration Fields

Field	Description
Task Group	The task group name.
Description	The associated description for task group name.
Parent Task Groups	The associated parent task groups for task group name. To configure this parent task group, click the Add icon in the header row. To remove the parent task group, click the Reset icon in the row.
Configured Permission	The configured task permissions for task group.
Configured Tasks	The list of task names. To configure this task, click the Add icon in the header row. To remove the task, click the Reset icon in the row. The tasks available are platform and package dependent.
Permissions	The task permissions. <ul style="list-style-type: none"> • Read • Write • Debug • Execute

Use the buttons to perform the following:

- To add a task group, click **Add** and specify a name for the group.
- To remove a task group, select the checkbox for the group you want to remove and click **Remove**.
- Click **Refresh** to update the information on the screen.

User Groups

The **User Group Configuration** page allows you to add, edit, and remove user groups.

To access this page, click **System > Users > User Groups** in the navigation menu.

Table 36: User Group Configuration Fields

Field	Description
User Group	The user group name.
Description	The associated description for the user group name.
Parent User Groups	The associated parent user groups for user group. To configure this parent user group, click the Add icon in the header row. To remove the parent user group, click the Reset icon in the row.
Configured Permission	The configured task permissions for task group.
Contained Task Group	The associated task groups for the user group. To configure the task group, click the Add icon in the header row. To remove the task group, click the Reset icon in the row.
Operational Permission	The operational task permissions for the user group. <ul style="list-style-type: none"> • Read • Write • Debug • Execute

Use the buttons to perform the following:

- To add a user group, click **Add** and specify a name for the group.
- To remove a user group, select the checkbox for the group you want to remove and click **Remove**.
- Click **Refresh** to update the information on the screen.

Accounting List Configuration: Accounting List

Use the **Accounting List Configuration** page to view and configure the accounting lists for users who access the CLI to manage and monitor the device. Accounting Lists are used to record user activity on the device. The device is preconfigured with accounting lists. These are default lists, and they cannot be deleted. Additionally, the List Name and Accounting Type settings for the default lists cannot be changed.

To access this page, click **System > AAA > Accounting List** in the navigation menu.

Use the buttons to perform the following tasks in the **Accounting List** tab:

- To configure a new accounting list, click **Add**.
- To edit a list, select the entry to modify and click **Edit**. The settings that can be edited depend on the list type.

- To remove a non-default accounting list, click the – (minus) button associated with the entry. You must confirm the action before the entry is deleted.
- To reset the Method Options for a default accounting list to the factory default values, click the **Reset** icon associated with the entry. You must confirm the action before the entry is reset.

After you click **Add** or **Edit**, a window opens and allows you to configure accounting list settings. When adding an accounting list, you can configure the List Name, Accounting Type, and Record Type fields as well as the Accounting Methods. When editing an existing authentication list, only the Record Type and Accounting Methods can be configured. The following information describes how to set the Accounting Methods.

Table 37: Add New Accounting List Fields

Field	Description
Accounting Methods	This area includes the Available Methods and Selected Methods fields. If a list uses multiple accounting methods, the order in which you move the method from the Available Methods field to the Selected Methods field determines the order in which the device attempts to send accounting notifications. If the device successfully sends the accounting notifications by using the first method, the next method is not attempted.
Available Methods	The accounting methods that can be used for the accounting list. To set the accounting method, select the method in the Available Methods field and click the right arrow to move it into the Selected Methods field.
Selected Methods	The accounting methods currently configured for the list. When multiple methods are in this field, the order in which the methods are listed is the order in which the methods will be used. If the device is unable to send accounting notifications by using the first method, the device attempts to send notifications by using the second method. To remove a method from this field, select it and click the left arrow to return it to the Available Methods area.

Table 38: Accounting List Configuration Fields

Field	Description
Accounting Type	The type of accounting list, which is one of the following: <ul style="list-style-type: none"> • Command: Each CLI command executed by the user, along with the time the command was executed, is recorded and sent to an external AAA server. • EXEC: User login and logout times are recorded and sent to an external AAA server.
List Name	The name of the accounting list. This field can be configured only when adding a new accounting list.
Record Type	Indicates when to record and send information about the user activity: <ul style="list-style-type: none"> • StartStop: Accounting notifications are sent at the beginning and at the end of an exec session or a user-executed command. User activity does not wait for the accounting notification to be recorded at the AAA server. • StopOnly: Accounting notifications are sent at the end of an exec session or a user-executed command.
Method Options	The method(s) used to record user activity. The possible methods are as follows: <ul style="list-style-type: none"> • TACACS+: Accounting notifications are sent to the configured <u>TACACS+</u> server. • RADIUS: Accounting notifications are sent to the configured <u>RADIUS</u> server.

Table 38: Accounting List Configuration Fields (continued)

Field	Description
List Type	The type of accounting list, which is one of the following: <ul style="list-style-type: none"> • Default: The list is preconfigured on the system. This type of list cannot be deleted, and only the Method Options and Record Type settings are configurable. • Configured: The list has been added by a user.
Access Line	The access method(s) that use the list for accounting user activity. The settings for this field are configured on the Accounting Selection page.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Accounting List Configuration: Accounting Selection

Use the **Accounting List Configuration** page to associate an accounting list with each access method. For each access method, the following two accounting lists are associated:

- Exec – The accounting list to record user login and logout times.
- Commands – The accounting list to record which actions a user takes on the system, such as page views or configuration changes. This list also records the time when the action occurred. For terminal access methods, this list records the CLI commands a user executes and when each command is issued.

To access this page, click **System > AAA > Accounting Selection** in the navigation menu.

Complete the following fields in the **Accounting Selection** tab:

Table 39: Accounting List Configuration Fields

Field	Description
Terminal	The access methods in this section are CLI-based. <ul style="list-style-type: none"> • Console: The Exec accounting list and the Commands accounting list to apply to users who access the CLI by using a connection to the console port. • Telnet: The Exec accounting list and the Commands accounting list to apply to users who access the CLI by using a <i>Telnet</i> session. • SSH: The Exec accounting list and the Commands accounting list to apply to users who access the CLI by using an <i>SSH</i> session.
Hypertext Transfer Protocol	The access methods in this section are through a web browser. <ul style="list-style-type: none"> • HTTP: The Exec accounting list and the Commands accounting list to apply to users who access the web-based management interface by using HTTP. • Telnet: The Exec accounting list and the Commands accounting list to apply to users who access the web-based management interface by using secure HTTP.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Authentication List Summary

Use the **Authentication List Summary** page to view and configure the authentication lists used for management access and port-based (IEEE 802.1X) access to the system. An authentication list specifies which authentication method(s) to use to validate the credentials of a user who attempts to access the device. Several authentication lists are preconfigured on the system. These are default lists, and they cannot be deleted. Additionally, the List Name and Access Type settings for the default lists cannot be changed.

To access this page, click **System > AAA > Authentication List** in the navigation menu.

Table 40 shows the fields for the Authentication List Summary page.

Table 40: Authentication List Summary Fields

Field	Description
List Name	The name of the authentication list. This field can be configured only when adding a new authentication list.
Access Type	<p>The way the user accesses the system. This field can be configured only when adding a new authentication list, and only the Login and Enable access types can be selected. The access types are as follows:</p> <ul style="list-style-type: none"> • Login: User EXEC-level management access to the command-line interface (CLI) by using a console connection or a <i>Telnet</i> or <i>SSH</i> session. Access at this level has a limited number of CLI commands available to view or configure the system. • Enable: Privileged EXEC-level management access to the CLI by using a console connection or a telnet or SSH session. In Privileged EXEC mode, read-write users have access to all CLI commands. • HTTP: Management-level access to the web-based user interface by using HTTP. • HTTPS: Management-level access to the web-based user interface by using secure HTTP. • Dot1x: Port-based access to the network through a switch port that is controlled by IEEE 802.1X.
Method Options	<p>The method(s) used to authenticate a user who attempts to access the management interface or network. The possible methods are as follows:</p> <ul style="list-style-type: none"> • Enable: Uses the locally configured Enable password to verify the user's credentials. • Line: Uses the locally configured Line password to verify the user's credentials. • Local: Uses the ID and password in the Local User database to verify the user's credentials. • RADIUS: Sends the user's ID and password to the configured <i>RADIUS</i> server to verify the user's credentials. • TACACS+: Sends the user's ID and password to the configured <i>superloop</i> server to verify the user's credentials. • None: No authentication is used. • IAS: Uses the local Internal Authentication Server (IAS) database for 802.1X port-based authentication.

Table 40: Authentication List Summary Fields (continued)

Field	Description
List Type	The type of list, which is one of the following: <ul style="list-style-type: none"> • Default: The list is preconfigured on the system. This type of list cannot be deleted, and only the Method Options are configurable. • Configured: The list has been added by a user.
Access Line	The access method(s) that use the list for authentication. The settings for this field are configured on the Authentication Selection page.

- Click **Refresh** to update the information on the screen.
- To create a new authentication list, see [Authentication Server Users](#) on page 59. To assign users to a specific authentication list, see [User Accounts](#) on page 57. To configure the 802.1x port security users, see [RADIUS Settings](#) on page 284.

Select Authentication List

Use the **Select Authentication List** Configuration page to associate an authentication list with each CLI-based access method (console, Telnet, and SSH). Each access method has the following two authentication lists associated with it:

- **Login** – The authentication list to use for User EXEC-level management access to the CLI. Access at this level has a limited number of CLI commands available to view or configure the system. The options available in this menu include the default Login authentication lists as well as any user-configured Login lists.
- **Enable** – The authentication list to use for Privileged EXEC-level management access to the CLI. In Privileged EXEC mode, read-write users have access to all CLI commands. The options available in this menu include the default Enable authentication lists as well as any user-configured Enable lists.

To access this page, click **System > AAA > Authentication Selection** in the navigation menu.

[Table 41](#) shows the fields for this page.

Table 41: Select Authentication List Fields

Field	Description
Console	The Login authentication list and the Enable authentication list to apply to users who attempt to access the CLI by using a connection to the console port.
Telnet	The Login authentication list and the Enable authentication list to apply to users who attempt to access the CLI by using a Telnet session.
Secure Telnet (SSH)	The Login authentication list and the Enable authentication list to apply to users who attempt to access the CLI by using a secure shell (SSH) session.
List Name	The name of the authentication list. This field can be configured only when adding a new authentication list.

Table 41: Select Authentication List Fields (continued)

Field	Description
Access Type	<p>The way the user accesses the system. This field can be configured only when adding a new authentication list, and only the Login and Enable access types can be selected. The access types are as follows:</p> <ul style="list-style-type: none"> • Login: User EXEC-level management access to the command-line interface (CLI) by using a console connection or a telnet or SSH session. Access at this level has a limited number of CLI commands available to view or configure the system. • Enable: Privileged EXEC-level management access to the CLI by using a console connection or a telnet or SSH session. In Privileged EXEC mode, read-write users have access to all CLI commands. • HTTP: Management-level access to the web-based user interface by using HTTP. • HTTPS: Management-level access to the web-based user interface by using secure HTTP. • Dot1x: Port-based access to the network through a switch port that is controlled by IEEE 802.1X.
Method Options	<p>The method(s) used to authenticate a user who attempts to access the management interface or network. The possible methods are as follows:</p> <ul style="list-style-type: none"> • Enable: Uses the locally configured Enable password to verify the user's credentials. • Line: Uses the locally configured Line password to verify the user's credentials. • Local: Uses the ID and password in the Local User database to verify the user's credentials. • RADIUS: Sends the user's ID and password to the configured <i>RADIUS</i> server to verify the user's credentials. • TACACS+: Sends the user's ID and password to the configured TACACS+ server to verify the user's credentials. • None: No authentication is used. • IAS: Uses the local Internal Authentication Server (IAS) database for 802.1X port-based authentication.
List Type	<p>The type of list, which is one of the following:</p> <ul style="list-style-type: none"> • Default: The list is preconfigured on the system. This type of list cannot be deleted, and only the Method Options are configurable. • Configured: The list has been added by a user.
Access Line	<p>The access method(s) that use the list for authentication. The settings for this field are configured on the Authentication Selection page.</p>

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Authorization List Configuration

Use the **Authorization List Configuration** page to view and configure the authorization lists for users who access the CLI and for users who access the network through IEEE 802.1X-enabled ports.

Authorization lists are used to determine whether a user is permitted to perform a given activity on the

system or network. Several authorization lists are preconfigured on the system. These are default lists, and they cannot be deleted. Additionally, the List Name and Authorization Type settings for the default lists cannot be changed.

To access this page, click **System > AAA > Authorization List** in the navigation menu.

Table 42: Authorization List Configuration Fields

Field	Description
List Name	The name of the authorization list. This field can be configured only when adding a new authorization list.
Authorization Type	The type of authorization list, which is one of the following: <ul style="list-style-type: none"> • Command: Determines which CLI commands a user is permitted to issue. When command authorization is enabled, each command a user enters must be validated before the command is executed. • EXEC: Determines whether a user can bypass User EXEC mode and enter Privileged EXEC mode directly after a successful Login authentication. • Network: Determines whether the user is permitted to access various network services. This authorization type applies to port-based access (IEEE 802.1X) rather than access to the CLI.
Method Options	The method(s) used to authorize a user's access to the device or network services. The possible methods are as follows: <ul style="list-style-type: none"> • TACACS+: When a user issues a CLI command, the device contacts the configured <u>TACACS+</u> server to verify whether the user is allowed to issue the command. If approved, the command is executed. Otherwise, the command fails. • RADIUS: When a user is authenticated by the <u>RADIUS</u> server, the device downloads a list of permitted/denied commands from the RADIUS server. The list of authorized commands that are associated with the authenticated user is cached during the user's session. If this method is selected, the authentication method for the access type must also be RADIUS. • Local: Uses a list stored locally on the system to determine whether the user is authorized to access the given services. • None: No authorization is used. If the method is None, the authorization type is effectively disabled.
List Type	The type of authorization list, which is one of the following: <ul style="list-style-type: none"> • Default: The list is preconfigured on the system. This type of list cannot be deleted, and only the Method Options are configurable. • Configured: The list has been added by a user.
Access Line	The access method(s) that use the list for authorization. The settings for this field are configured on the Authorization Selection page.

- To configure a new authorization list, click **Add**.
- To edit a list, select the entry to modify and click **Edit**. The settings that can be edited depend on the list type.
- To remove a non-default authorization list, click the – (minus) button associated with the entry. You must confirm the action before the entry is deleted.

To reset the Method Options for a default authorization list to the factory default values, click the **Reset** icon associated with the entry. You must confirm the action before the entry is reset.

After you click **Add** or **Edit**, a window opens and allows you to configure authorization list settings.

When adding an authorization list, you can configure the List Name and Authorization Type fields as well as the Authorization Methods. When editing an existing authentication list, only the Authorization Methods can be configured. The following information describes how to set the Authorization Methods.

Table 43: Add New Authorization List Fields

Field	Description
Authorization Methods	This area includes the Available Methods and Selected Methods fields. For lists that allow multiple authorization methods, the order in which you move the method from the Available Methods field to the Selected Methods field determines the order in which the device attempts to authorize the user.
Available Methods	The authorization methods that can be used for the authorization list. Not all methods are available for all lists. To set the authorization method, select the method in the Available Methods field and click the right arrow to move it into the Selected Methods field.
Selected Methods	The authorization methods currently configured for the list. When multiple methods are in this field, the order in which the methods are listed is the order in which the methods will be used to authorize a user. If the user fails to be authorized using the first method, the device attempts to authorize the user by using the next method in the list. No authorization methods can be added after None. To remove a method from this field, select it and click the left arrow to return it to the Available Methods area.

Line Password

Use the **Line Password** page to configure line mode passwords.

To display this page, click **System > Passwords > Line Password** in the navigation menu.

Table 44: Line Password Fields

Field	Description
Line Mode	Any or all of the following passwords may be changed on this page by checking the box that precedes it: <ul style="list-style-type: none"> • Console • Telnet • SSH
Password (8-64 characters)	Enter the new password for the corresponding Line Mode in this field. Be sure the password conforms to the allowed number of characters. The password characters are not displayed on the page, but are disguised in a browser-specific manner.
Confirm Password (8-64 characters)	Re-enter the new password for the corresponding Line Mode in this field. This must be the same value entered in the Password field. Be sure the password conforms to the allowed number of characters. The password characters are not displayed on the page, but are disguised in a browser-specific manner.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Enable Password

Use the **Enable Password** page to configure the enable password.

To display this page, click **System > Passwords > Enable Password** in the navigation menu.

Table 45: Enable Password Fields

Field	Description
Enable Password	Specify the password all users must enter after executing the enable command at the CLI prompt.
Confirm Enable Password	Confirms the new enable password. The password appears in the **** format.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Password Rules

Use the **Password Rules** page to configure settings that apply to all user passwords.

To display this page, click **System > Passwords > Password Rules** in the navigation menu.

Table 46: Password Rules Configuration Fields

Field	Description
Minimum Length	Passwords must have at least this many characters (range is 8 to 64).
Aging (days)	The number of days that a user password is valid from the time the password is set. Once a password expires, the user is required to enter a new password at the next login.
History	The number of previous passwords that are retained to prevent password reuse. This helps to ensure that a user does not attempt to reuse the same password too often.
Lockout Attempts	After a user fails to log in this number of times, the user is locked out until the password is reset by the administrator.
Strength Check	Enable or disable the password strength check feature. Enabling this feature forces the user to configure passwords that comply with the strong password configuration specified in the following fields.
Minimum Number of Uppercase Letters	Specify the minimum number of uppercase letters a password must include.
Minimum Number of Lowercase Letters	Specify the minimum number of lowercase letters a password must include.
Minimum Number of Numeric Characters	Specify the minimum number of numbers a password must include.
Minimum Number of Special Characters	Specify the minimum number of special characters (non-alphanumeric, such as # or &) a password must include.
Maximum Number of Repeated Characters	Specify the maximum number of repeated characters a password is allowed to include. An example of four repeated characters is aaaa.

Table 46: Password Rules Configuration Fields (continued)

Field	Description
Maximum Number of Consecutive Characters	Specify the maximum number of consecutive characters a password is allowed to include. An example of four consecutive characters is abcd.
Minimum Character Classes	Specify the minimum number of character classes a password must contain. There are four character classes: <ul style="list-style-type: none"> • Uppercase • Lowercase • Numbers • Special Characters
Exclude Keyword	The password to be configured should not contain the keyword mentioned in this field. The valid range for the keyword is 2 to 64 characters.
Exclude Keyword Name	The list of keywords that a valid password must not contain. Excluded keyword checking is case-insensitive. Additionally, a password cannot contain the backwards version of an excluded keyword. For example, if pass is an excluded keyword, passwords such as 23passA2c, ssapword, and PAsSwoRD are prohibited. Use the plus and minus buttons to perform the following tasks: <ul style="list-style-type: none"> • To add a keyword to the list, click the + (plus) button, type the word to exclude in the Exclude Keyword Name field, and click Submit. • To remove a keyword from the list, click the - (minus) button associated with the keyword to remove and confirm the action. • To remove all keywords from the list, click the - (minus) button in the header row and confirm the action.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Last Password Result

Use the **Last Password Result** page view information about the last attempt to set a user password. If the password set was unsuccessful, a reason for the failure is given.

To display this page, click **System > Password > Last Password** in the navigation menu.

Table 47: Last Password Result

Field	Description
Last Password Set Result	Displays information about the last (User/Line/Enable) password configuration result. If the field is blank, no passwords have been configured on the device. Otherwise, the field shows that the password was successfully set or provides information about the type of password configuration that failed and why it could not be set.
Strength Check	Displays Enabled if Strength Check is applied in last password change, otherwise it displays Disabled.

Denial of Service

Use the **Denial of Service** (DoS) page to configure DoS control. 200 Series software provides support for classifying and blocking specific types of DoS attacks. You can configure your system to monitor and block these types of attacks:

- SIP=DIP: Source IP address = Destination IP address.
- First Fragment: TCP Header size smaller than configured value.
- TCP Fragment: IP Fragment Offset = 1.
- TCP Flag: TCP Flag SYN set and Source Port < 1024 or TCP Control Flags = 0 and TCP Sequence Number = 0 or TCP Flags FIN, URG, and PSH set and TCP Sequence Number = 0 or TCP Flags SYN and FIN set.
- L4 Port: Source TCP/UDP Port = Destination TCP/UDP Port.
- ICMP: Limiting the size of *ICMP (Internet Control Message Protocol)* Ping packets.

To access this page, click **System > Advanced Configuration > Protection > Denial of Service** in the navigation menu.

Table 48: Denial of Service Configuration Fields

Field	Description
TCP Settings	
First Fragment	Enable this option to allow the device to drop packets that have a TCP header smaller than the value configured in the Min TCP Hdr Size field.
TCP Port	Enable this option to allow the device to drop packets that have the TCP source port equal to the TCP destination port.
UDP Port	Enable this option to allow the device to drop packets that have the UDP source port equal to the UDP destination port.
SIP=DIP	Enable this option to allow the device to drop packets that have a source IP address equal to the destination IP address.
SMAC=DMAC	Enable this option to allow the device to drop packets that have a source MAC address equal to the destination MAC address.
TCP FIN and URG and PSH	Enable this option to allow the device to drop packets that have TCP Flags FIN, URG, and PSH set and a TCP Sequence Number equal to 0.
TCP Flag and Sequence	Enable this option to allow the device to drop packets that have TCP control flags set to 0 and the TCP sequence number set to 0.
TCP SYN	Enable this option to allow the device to drop packets that have TCP Flags SYN set.
TCP SYN and FIN	Enable this option to allow the device to drop packets that have TCP Flags SYN and FIN set.
TCP Fragment	Enable this option to allow the device to drop packets that have a TCP payload where the IP payload length minus the IP header size is less than the minimum allowed TCP header size.
TCP Offset	Enable this option to allow the device to drop packets that have a TCP header Offset set to 1.

Table 48: Denial of Service Configuration Fields (continued)

Field	Description
Port D-Disable	Enable this option to allow the system to diagnostically disable an interface if a potential DoS attack has been detected on that interface. If an interface is diagnostically disabled, it remains in the disabled state until an administrator manually enables the interface.
Min TCP Hdr Size	The minimum TCP header size allowed. If First Fragment DoS prevention is enabled, the device will drop packets that have a TCP header smaller than this configured value.
ICMP Settings: These options help prevent the device and the network from attacks that involve issues with the ICMP echo request packets (pings) that the device receives.	
ICMP	Enable this option to allow the device to drop ICMP packets that have a type set to ECHO_REQ (ping) and a payload size greater than the ICMP payload size configured in the Max ICMPv4 Size or Max ICMPv6 Size fields.
ICMP Fragment	Enable this option to allow the device to drop fragmented ICMP packets.
Max ICMPv4 Size	The maximum allowed ICMPv4 packet size. If ICMP DoS prevention is enabled, the device will drop ICMPv4 ping packets that have a size greater than this configured maximum ICMPv4 packet size.
Max ICMPv6 Size	The maximum allowed IPv6 ICMP packet size. If ICMP DoS prevention is enabled, the switch will drop IPv6 ICMP ping packets that have a size greater than this configured maximum ICMPv6 packet size.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Managing the DHCP Server

[DHCP](#) is generally used between clients (for example, hosts) and servers (for example, routers) for the purpose of assigning IP addresses, gateways, and other networking definitions such as DNS, NTP, and/or SIP parameters. The DHCP Server folder contains links to web pages that define and display DHCP parameters and data.

Global Configuration

Use the **Global Configuration** page to configure global parameters for [DHCP](#).

To display this page, click **System > Advanced Configuration > DHCP Server > Global** in the navigation menu.

Table 49: DHCP Server Global Configuration Fields

Field	Description
Admin Mode	Enables or disables the DHCP server administrative mode. When enabled, the device can be configured to automatically allocate TCP/IP configurations for clients.
Conflict Logging Mode	Enables or disables the logging mode for IP address conflicts. When enabled, the system stores information IP address conflicts that are detected by the DHCP server.

Table 49: DHCP Server Global Configuration Fields (continued)

Field	Description
Bootp Automatic Mode	Enables or disables the BOOTP automatic mode. When enabled, the DHCP server supports the allocation of automatic addresses for BOOTP clients. When disabled the DHCP server supports only static addresses for BOOTP clients.
Ping Packet Count	The number of packets the server sends to a pool address to check for duplication as part of a ping operation. If the server receives a response to the ping, the address is considered to be in conflict and is removed from the pool.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Click **Refresh** to refresh the page with the most current data.

DHCP Server Excluded Addresses

Use the **DHCP Server Excluded Addresses** page to view and configure the IP addresses the *DHCP* server should not assign to clients

To display this page, click **System > Advanced Configuration > DHCP Server > Excluded Addresses** in the navigation menu.

Use the buttons to perform the following tasks:

- To add one or more IP addresses to exclude, click **Add** and specify the IPv4 address or range of addresses in the available fields. Then click **Submit**.
- To remove an excluded address or range of addresses, select each entry to remove and click **Remove**. You must confirm the action before the entries are removed.

Table 50: DHCP Server Excluded Addresses Fields

Field	Description
From	The IP address to exclude. In a range of addresses, this value is the lowest address to exclude.
To	The highest address to exclude in a range of addresses. If the excluded address is not part of a range, this field shows the same value as the From field. When adding a single IP address to exclude, you can enter the same address specified in the From field or leave the field with the default value.

Click **Refresh** to refresh the page with the most current data.

Pool Summary

Use the **DHCP Pool Summary** page to view the currently configured *DHCP* server pools and to add and remove pools. A DHCP server pool is a set of network configuration information available to DHCP clients that request the information..

To access this page, click **System > Advanced Configuration > DHCP Server > Pool Summary** in the navigation menu.

Use the buttons to perform the following tasks:

- To add a pool, click **Add** and configure the pool information in the available fields. Then click **Submit**.
- To remove one or more pools, select each entry to remove and click **Remove**. You must confirm the action before the pools are removed.

Table 51: DHCP Pool Summary Fields

Field	Description
Name	The name that identifies the DHCP server pool.
Type of Binding	The type of binding for the pool. The options are: <ul style="list-style-type: none"> • Manual : The DHCP server assigns a specific IP address to the client based on the client's MAC address. This type is also known as Static. • Dynamic: The DHCP server can assign the client any available IP address within the pool. This type is also known as Automatic. • Undefined: The pool has been created by using the CLI (command-line interface), but the pool information has not been configured.
Network	For a Manual pool, the host IP address to assign the client. For a Dynamic pool, the network base address.
Lease Time	The amount of time the information the DHCP server allocates is valid.
When you click Add , the Add DHCP Server Pool page opens and allows you to configure the following DHCP pool settings:	
Name	The name that identifies the DHCP server pool.
Type of Binding	The type of binding for the pool. The options are: <ul style="list-style-type: none"> • Manual: The DHCP server assigns a specific IP address to the client based on the client's MAC address. • Dynamic: The DHCP server can assign the client any available IP address within the pool. <p>The binding type you select determines the fields that are available to configure.</p>
Network Base Address	(Dynamic pools only) The network portion of the IP address. A DHCP client can be offered any available IP address within the defined network as long as it has not been configured as an excluded address.
Network Mask	(Dynamic pools only) The subnet mask associated with the Network Base Address that separates the network bits from the host bits.
Client Name (optional)	(Manual pools only) The system name of the client. The Client Name should not include the domain name.
Hardware Address Type	(Manual pools only) The protocol type (Ethernet or IEEE 802) used by the client's hardware platform. This value is used in response to requests from BOOTP clients.
Hardware Address	(Manual pools only) The MAC address of the client.
Client ID	(Manual pools only) The value some DHCP clients send in the Client Identifier field of DHCP messages. This value is typically identical to the Hardware Address value. In some systems, such as Microsoft DHCP clients, the client identifier is required instead of the hardware address. If the client's DHCP request includes the client identifier, the Client ID field on the DHCP server must contain the same value, and the Hardware Address Type field must be set to the appropriate value. Otherwise, the DHCP server will not respond to the client's request.
Host IP Address	(Manual pools only) The IP address to offer the client.

Table 51: DHCP Pool Summary Fields (continued)

Field	Description
Host Mask	(Manual pools only) The subnet mask to offer the client.
Lease Expiration Mode	Whether the information the server provides to the client should expire. The options are: <ul style="list-style-type: none"> • Enable: Allows the lease to expire. If you select this option, you can specify the amount of time the lease is valid in the Lease Duration field. • Disable: Sets an infinite lease time. For dynamic bindings, an infinite lease time implies a lease period of 60 days. For a manual binding, an infinite lease period never expires.
Lease Duration	The number of Days, Hours, and Minutes the lease is valid. This field cannot be configured if the Lease Expiration Mode is disabled.
Default Router Address (optional)	The IP address of the router to which the client should send traffic. The default router should be in the same subnet as the client. To add additional default routers, use the DHCP Server Pool Configuration page.
DNS Server Address (optional)	The IP addresses of up to two DNS servers the client should use to resolve host names into IP addresses. To add additional DNS servers, use the DHCP Server Pool Configuration page.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Pool Configuration

Use the **DHCP Pool Configuration** page to edit pool settings or to configure additional settings for existing manual and dynamic pools. The fields that can be configured depend on the type of binding that is selected.

To access this page, click **System > Advanced Configuration > DHCP Server > Pool Configuration** in the navigation menu.

If you select **Automatic** or **Manual** from the **Type of Binding** drop-down menu, the screen refreshes and a slightly different set of fields appears.

Table 52: DHCP Pool Configuration Fields

Field	Description
Pool Name	For a user with read/write permission, this field would show names of all the existing pools along with an additional option Create. When the user selects Create, another text box, Pool Name, appears where the user may enter name for the Pool to be created. For a user with read-only permission, this field would show names of the existing pools only.
Type of Binding	Specifies the type of binding for the pool. <ul style="list-style-type: none"> • Unallocated: The addresses are not assigned to a client. • Automatic: The IP address is automatically assigned to a client by the DHCP server. • Manual: You statically assign an IP address to a client based on the client's MAC address.

Table 52: DHCP Pool Configuration Fields (continued)

Field	Description
Network Base Address	(Dynamic pools only) The network portion of the IP address. A DHCP client can be offered any available IP address within the defined network as long as it has not been configured as an excluded address.
Network Mask	For dynamic bindings, this field specifies the subnet mask for a DHCP address of a dynamic pool. You can enter a value in Network Mask or Prefix Length to specify the subnet mask, but do not enter a value in both fields.
Client Name	For manual bindings, this field specifies a name for the client to which the DHCP server will statically assign an IP address. This field is optional.
Hardware Address Type	For manual bindings, this field specifies the protocol of the hardware platform of the DHCP client. Valid types are ethernet and ieee802. Default value is ethernet.
Hardware Address	For manual bindings, this field specifies the MAC address of the hardware platform of the DHCP client.
Client ID	(Manual pools only) The value some DHCP clients send in the Client Identifier field of DHCP messages. This value is typically identical to the Hardware Address value. In some systems, such as Microsoft DHCP clients, the client identifier is required instead of the hardware address. If the client's DHCP request includes the client identifier, the Client ID field on the DHCP server must contain the same value, and the Hardware Address Type field must be set to the appropriate value. Otherwise, the DHCP server will not respond to the client's request.
Host IP Address	(Manual pools only) The IP address to offer the client.
Host Mask	For manual bindings, this field specifies the subnet mask to be statically assigned to a DHCP client. You can enter a value in Host Mask or Prefix Length to specify the subnet mask, but do not enter a value in both fields.
Lease Expiration	Whether the information the server provides to the client should expire. <ul style="list-style-type: none"> • Enable: Allows the lease to expire. If you select this option, you can specify the amount of time the lease is valid in the Lease Duration field. • Disable: Sets an infinite lease time. For Dynamic bindings, an infinite lease time implies a lease period of 60 days. For a Manual binding, an infinite lease period never expires.
Lease Duration	The number of Days, Hours, and Minutes the lease is valid. This field cannot be configured if the Lease Expiration is disabled.
Next Server Address	The IP address of the next server the client should contact in the boot process. For example, the client might be required to contact a TFTP server to download a new image file. To configure this field, click the Edit icon in the row. To reset the field to the default value, click the Reset icon in the row. To configure settings for one or more default routers, DNS servers, or NetBIOS servers that can be used by the client(s) in the pool, use the buttons available in the appropriate table to perform the following tasks: <ul style="list-style-type: none"> • To add an entry to the server list, click the + (plus) button and enter the IP address of the server to add. • To edit the address of a configured server, click the Edit icon associated with the entry to edit and update the address. • To delete an entry from the list, click the - (minus) button associated with the entry to remove. • To delete all entries from the list, click the - (minus) button in the heading row.

Table 52: DHCP Pool Configuration Fields (continued)

Field	Description
Default Router	Lists the IP address of each router to which the client(s) in the pool should send traffic. The default router should be in the same subnet as the client.
DNS Server	Lists the IP address of each DNS server the client(s) in the pool can contact to perform address resolution.
NetBIOS Server	Lists the IP address of each NetBIOS Windows Internet Naming Service (WINS) name server that is available for the selected pool.

After you configure values for the DHCP address pool, click **Submit** to create the pool and apply the changes to the system.

To delete a pool, select the pool from the **Pool Name** drop-down menu and click **Delete**.

Pool Options

Use the **Pool Options** page to configure additional *DHCP* pool options, including vendor-defined options. DHCP options are collections of data with type codes that indicate how the options should be used. When a client broadcasts a request for information, the request includes the option codes that correspond to the information the client wants the DHCP server to supply.

To access this page, click **System > Advanced Configuration > DHCP Server > Pool Options** in the navigation menu.

If no DHCP pools exist, the **Pool Options** page does not display the fields shown in [Table 53](#).

If any DHCP pools are configured on the system, this page contains the following fields:

Table 53: Pool Options Fields

Field	Description
Pool Name	Select the DHCP pool to with the options you want to view or configure.
Option Code	Displays the DHCP option code configured for the selected Pool.
Option Type	Specifies the type of option associated with the option code configured for the selected pool. The possible values are as follows: <ul style="list-style-type: none"> • Ascii: The option type is a text string. • Hex: The option type is a hexadecimal number. • IP Address: The option type is an IP address.
ASCII Value	Shows the Option ASCII Value for the selected pool.
Hex Value	Shows the Option Hex Value for the selected pool.
IP Address Value	Shows the Option IP Address Value for the selected pool.
Delete Option Code	To delete an option code for the selected Pool, enter the option code in the folder and click Delete . This button is not visible to a user with read-only permission.

Bindings Information

Use the **Bindings Information** page to view information about the IP address bindings in the *DHCP* server database.

To access this page, click **System > Advanced Configuration > DHCP Server > Bindings** in the navigation menu.

Table 54: Bindings Information Fields

Field	Description
IP Address	The IP Address of the DHCP client.
Hardware Address	The MAC address of the DHCP client.
Lease Time Left	The amount of time left until the lease expires in days, hours, and minutes.
Pool Allocation Type	The type of binding used: <ul style="list-style-type: none"> • Dynamic: The address was allocated dynamically from a pool that includes a range of IP addresses. • Manual: A static IP address was assigned based on the MAC address of the client. • Inactive: The pool is not in use.
Clear Entries (Button)	To remove an entry from the table, select each entry to delete and click Clear Entries . You must confirm the action before the binding is deleted.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Server Statistics

Use the **DHCP Server Statistics** page to view information about the DHCP server bindings and messages.

To access this page, click **System > Advanced Configuration > DHCP Server > Statistics** in the navigation menu.

Table 55: Server Statistics Fields

Field	Description
Automatic Bindings	Shows the number of automatic bindings on the DHCP server.
Expired Bindings	Shows the number of expired bindings on the DHCP server.
Malformed Messages	Shows the number of the malformed messages.
Message Received	
DHCPDISCOVER	Shows the number of DHCPDISCOVER messages received by the DHCP server.
DHCPREQUEST	Shows the number of DHCPREQUEST messages received by the DHCP server.
DHCPDECLINE	Shows the number of DHCPDECLINE messages received by the DHCP server.
DHCPRELEASE	Shows the number of DHCPRELEASE messages received by the DHCP server.
DHCPINFORM	Shows the number of DHCPINFORM messages received by the DHCP server.

Table 55: Server Statistics Fields (continued)

Field	Description
DHCPOFFER	Shows the number of DHCPOFFER messages sent by the DHCP server.
DHCPACK	Shows the number of DHCPACK messages sent by the DHCP server.
DHCPNAK	Shows the number of DHCPNAK messages sent by the DHCP server.
Message Sent	
DHCPOFFER	The number of DHCP offer messages the DHCP server has sent to DHCP clients in response to DHCP discovery messages it has received.
DHCPACK	The number of DHCP acknowledgement messages the DHCP server has sent to DHCP clients in response to DHCP request messages it has received. The server sends this message after the client has accepted the offer from this particular server. The DHCP acknowledgement message includes information about the lease time and any other configuration information that the DHCP client has requested.
DHCPNAK	The number of negative DHCP acknowledgement messages the DHCP server has sent to DHCP clients. A server might send this type of message if the client requests an IP address that is already in use or if the server refuses to renew the lease.

- Click **Refresh** to update the information on the screen.
- Click **Clear Server Statistics** to reset all counters to zero.

Conflicts Information

Use the **Conflicts Information** page to view information on hosts that have address conflicts; that is, when the same IP address is assigned to two or more devices on the network.

To access this page, click **System > Advanced Configuration > DHCP Server > Conflicts** in the navigation menu.

Table 56: Conflicts Information Fields

Field	Description
IP Address	The IP address that has been detected as a duplicate.
Detection Method	The method used to detect the conflict, which is one of the following: <ul style="list-style-type: none"> • Gratuitous ARP: The <i>DHCP</i> client detected the conflict by broadcasting an ARP request to the address specified in the DHCP offer message sent by the server. If the client receives a reply to the ARP request, it declines the offer and reports the conflict. • Ping: The server detected the conflict by sending an <i>ICMP</i> echo message (ping) to the IP address before offering it to the DHCP client. If the server receives a response to the ping, the address is considered to be in conflict and is removed from the pool. • Host Declined: The server received a DHCPDECLINE message from the host. A DHCPDECLINE message indicates that the host has discovered that the IP address is already in use on the network.

Table 56: Conflicts Information Fields (continued)

Field	Description
Detection Time	The time when the conflict was detected in days, hours, minutes and seconds since the system was last rebooted (that is, system up time).
Clear Entries (Button)	Clears all of the address conflict entries.

Configuring DNS

You can use these pages to configure information about *DNS (Domain Name Server)* servers the network uses and how the switch/router operates as a DNS client.

Global Configuration

Use the **Configuration** page to configure global DNS settings and to view DNS client status information.

To access this page, click **System > Advanced Configuration > DNS > Configuration**.

Table 57: DNS Global Configuration Fields

Field	Description
Admin Mode	Select Enable or Disable from the drop-down menu to set the administrative status of DNS Client. The default is Disable.
Default Domain Name	Enter the default domain name for DNS client messages. The name should be no longer than 255 characters. When the system is performing a lookup on an unqualified hostname, this field is provided as the domain name (for example, if the default domain name is ".com" and the user enters hotmail, then hotmail is changed to hotmail.com to resolve the name). By default, no default domain name is configured in the system.
Retry Number	Enter the number of times to retry sending DNS queries. The valid values are from 0 to 100. The default value is 2.
Response Timeout	Enter the number of seconds to allow a DNS server to respond to a request before issuing a retry. Valid values are 0 to 3600. The default value is 3.
Domain List	Enter a domain list to define the domain to use when performing a lookup on an unqualified hostname. Each name must be no more than 256 characters. Multiple default domain names can be configured using the default domain-name list. If there is no domain list, the default domain name configured is used.

- If you change any settings, click **Submit** to send the information to the system.
- To create a new list of domain names, click **Create**. Then enter a name of the list and click submit. Repeat this step to add multiple domains to the default domain list.
- To remove a domain from the default list select the **Remove** option next to the item you want to remove and click **Submit**.

DNS Host Name IP Mapping Configuration

Use the **IP Mapping** page to configure DNS host names for hosts on the network and to view dynamic DNS entries. The host names are associated with IPv4 or IPv6 addresses on the network, which are statically assigned to particular hosts.

To access this page, click **System > Advanced Configuration > DNS > IP Mapping** in the menu.

Table 58: DNS Host Name IP Mapping Summary Fields

Field	Description
DNS Static Entries	
Entry Type	Type of DNS entry: <ul style="list-style-type: none"> • Static: An entry that has been manually configured on the device. • Dynamic: An entry that the device has learned by using a configured DNS server to resolve a hostname.
Host Name	The name that identifies the system. For Static entries, specify the Host Name after you click Add .
IP Address	The IPv4 or IPv6 address associated with the configured Host Name. For Static entries, specify the IP Address after you click Add . You can specify either an IPv4 or an IPv6 address.
DNS Dynamic Entries	
Total Time	The number of seconds that the entry will remain in the table.
Elapsed Time	The number of seconds that have passed since the entry was added to the table. When the Elapsed Time reaches the Total Time, the entry times out and is removed from the table.
Dynamic Type	The type of address in the entry, for example IP or (less common) X.121.

The page includes the following command buttons:

- Click **Add Static Entry** to load the Host Name IP Mapping Configuration page in order to configure the Host Name IP Mapping entries.
- Click **Submit** to apply the new configuration and cause the change to take effect immediately. These changes will not be retained across a power cycle unless a Save is performed.
- Click **Clear Dynamic Entries** to remove all Host Name IP Mapping entries. A confirmation prompt will be displayed. Click the button to confirm removal and the Host Name IP Mapping dynamic entries are cleared.
- Click Refresh to refresh the page with the most current data from the switch.

If you click **Add**, the DNS Host Name IP Mapping Configuration window opens.

Table 59: DNS Host Name Mapping Configuration Fields

Field	Description
Host Name	Enter the host name to assign to the static entry.
IP Address	Enter the IP4 or IPv6 address associated with the host name.

The page includes the following command buttons:

- Click **Submit** to apply the new configuration and cause the change to take effect immediately. These changes will not be retained across a power cycle unless a Save is performed.
- Click **Cancel** to cancel and redisplay the hostname IP mapping page to see the configured hostname-IP mapping entries.

DNS Source Interface Configuration

Use the **DNS Source Configuration** page to specify the physical or logical interface to use as the DNS client source interface. When an IP address is configured on the source interface, this address is used for all DNS communications between the local DNS client and the remote DNS server. The IP address of the designated source interface is used in the IP header of DNS management protocol packets. This allows security devices, such as firewalls, to identify all source packets coming from a specific device.

To access the **DNS Source Interface Configuration** page, click **System > Advanced Configuration > DNS > Source Interface Configuration** in the menu.

Table 60: DNS Source Interface Configuration Fields

Field	Description
Type	The type of interface to use as the source interface: <ul style="list-style-type: none"> • None: The primary IP address of the originating (outbound) interface is used as the source address. • Interface: The primary IP address of a physical port is used as the source address. • Loopback: The primary IP address of the loopback interface is used as the source address. A loopback is always reachable, as long as any routing interface is up. • VLAN: The primary IP address of a VLAN routing interface is used as the source address.
Interface	When the selected Type is Interface, select the physical port to use as the source interface.
Loopback Interface	When the selected Type is Loopback, select the loopback interface to use as the source interface.
VLAN ID	When the selected Type is VLAN, select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.
Tunnel ID	When the selected Type is Tunnel, select the tunnel interface to use as the source interface.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Configuring Email Alerts

With the email alerting feature, log messages can be sent to one or more email addresses. You must configure information about the network *SMTP (Simple Mail Transfer Protocol)* server for email to be successfully sent from the switch.

The pages available from the Email Alerting folder allow you to configure information about what type of log message are sent via email and to what address(es) the messages are emailed.

Email Alert Global Configuration

Use the **Email Alert Global Configuration** page to configure the common settings for log messages emailed by the switch.

To access this page, click **System > Advanced Configuration > Email Alerts > Global** in the navigation menu.

Table 61: Email Alert Global Configuration Fields

Field	Description
Admin Mode	Sets the administrative mode of the feature. <ul style="list-style-type: none"> • Enable: The device can send email alerts to the configured SMTP server. • Disable: The device will not send email alerts.
From Address	Specifies the email address of the sender (the switch).
Log Duration	This duration in minutes determines how frequently the non critical messages are sent to the SMTP Server.
Urgent Messages Severity	Configures the severity level for log messages that are considered to be urgent. Messages in this category are sent immediately. The security level you select and all higher levels are urgent: <ul style="list-style-type: none"> • emergency (0): The device is unusable. • alert (1): Action must be taken immediately. • critical (2): The device is experiencing primary system failures. • error (3): The device is experiencing non-urgent failures. • warning (4): The device is experiencing conditions that could lead to system errors if no action is taken. • notice (5): The device is experiencing normal but significant conditions. • info (6): The device is providing non-critical information. • debug (7): The device is providing debug-level information.
Non Urgent Messages Severity	Configures the severity level for log messages that are considered to be nonurgent. Messages in this category are collected and sent in a digest form at the time interval specified by the Log Duration field. The security level you select and all levels up to, but not including the lowest Urgent level are considered nonurgent. Messages below the security level you specify are not sent via email. See the Urgent Message field description for information about the security levels.
Traps Severity	Configures the severity level for trap log messages. See the Urgent Message field description for information about the security levels.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

After configuring all email alert settings, click **Test** to send a test message to the configured address(es).

Email Alert Server Configuration

Use the **Email Alert Server Configuration** page to configure information about up to three SMTP (mail) servers on the network that can handle email alerts sent from the switch.

To access this page, click **System > Advanced Configuration > Email Alerts > Server** in the navigation menu.

Use the buttons to perform the following tasks:

- To add an SMTP server, click **Add** and configure the desired settings.
- To change information for an existing SMTP server, select the checkbox for the entry and click **Edit**. You cannot edit the host name or address of a server that has been added.
- To delete a configured SMTP server from the list, select the checkbox for the entry to delete and click **Remove**.

Table 62: Email Alert Server Configuration Fields

Field	Description
Host Name or IP Address	Shows the address or host name of the SMTP server that handles email alerts that the device sends.
Security	Specifies the type of authentication to use with the mail server, which can be TLSv1 (SMTP over SSL) or None (no authentication is required).
Port	Specifies the TCP port that email alerts are sent to on the SMTP server.
User Name	If the Security is TLSv1, this field specifies the user name required to access the mail server.
Password	If the Security is TLSv1, this field specifies the password associated with the configured user name for mail server access. When adding or editing the server, you must retype the password to confirm that it is entered correctly.

To remove a configured SMTP server, select the **Remove** checkbox and click **Delete**.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Email Alert Statistics

Use the **Email Alert Statistics** page to view information about email alerts sent from the switch.

To access this page, click **System > Advanced Configuration > Email Alerts > Statistics** in the navigation menu.

Table 63: Email Alert Statistics Fields

Field	Description
Number of Emails Sent	The number of email alert messages sent since the system was last rebooted.
Number of Emails Failed	The number of email alert messages that could not be sent since the system was last rebooted.
Time Since Last Email Sent	The amount of time that has passed since the last email alert message was sent successfully.

To update the page with the most current information, click **Refresh**. To reset the values on the page to zero, click **Clear Counters**.

Email Alert Subject Configuration

Use the **Email Alert Subject Configuration** page to configure the subject line of the email alert messages sent from the switch.

To access this page, click **System > Advanced Configuration > Email Alerts > Subject** in the navigation menu.

Table 64: Email Alert Subject Configuration Fields

Field	Description
Message Type	Select the appropriate option to configure the subject line of Urgent messages or Nonurgent messages.
Email Subject	Specify the text to be displayed in the subject of the email alert message.
Remove	To reset the email alert subject to the default value, select the Remove option associated with the message type to reset, and click Delete .

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Email Alert To Address Configuration

Use the **Email Alert To Address Configuration** page to configure the email addresses to which alert messages sent.

To access this page, click **System > Advanced Configuration > Email Alerts > Address** in the navigation menu.

Use the buttons to perform the following tasks:

- To add an email address to the list of email alert message recipients, click **Add** and configure the desired settings.
- To delete an entry from the list, select the checkbox for each entry to delete and click **Remove**.

Table 65: Email Alert To Address Configuration Fields

Field	Description
Message Type	Select the appropriate option to configure email address where Urgent messages or Nonurgent messages are sent.
To Address	Specify the email address to which the selected type of messages are sent.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Configuring and Viewing ISDP Information

The Industry Standard Discovery Protocol (ISDP) is a proprietary Layer 2 network protocol which inter-operates with Cisco devices running the *CDP (Cisco Discovery Protocol)*. ISDP is used to share information between neighboring devices. 200 Series software participates in the CDP protocol and is able to both discover and be discovered by other CDP supporting devices.

The following pages are accessible from this ISDP folder:

- [ISDP Global Configuration](#) on page 87
- [ISDP Cache Table](#) on page 87
- [ISDP Interface Configuration](#) on page 88
- [ISDP Statistics](#) on page 89

ISDP Global Configuration

To access the **ISDP Global Configuration** page, click **System > Advanced Configuration > ISDP > Global** in the navigation menu.

The following table describes the fields available on the page.

Table 66: ISDP Global Configuration Fields

Field	Description
ISDP Mode	Use this field to enable or disable the Industry Standard Discovery Protocol on the switch.
ISDP V2 Mode	Use this field to enable or disable the Industry Standard Discovery Protocol v2 on the switch.
Message Interval	Specifies the ISDP transmit interval. The range is (5–254). Default value is 30 seconds.
Hold Time Interval	The receiving device holds ISDP message during this time period. The range is (10–255). Default value is 180 seconds.
Device ID	The Device ID advertised by this device. The format of this Device ID is characterized by the value of Device ID Format object.
Device ID Format Capability	Indicates the Device ID format capability of the device. <ul style="list-style-type: none"> • serialNumber: Indicates that the device uses serial number as the format for its Device ID. • macAddress: Indicates that the device uses Layer 2 MAC address as the format for its Device ID. • other: Indicates that the device uses its platform specific format as the format for its Device ID.
Device ID Format	Indicates the Device ID format of the device. <ul style="list-style-type: none"> • serialNumber: Indicates that the value is in the form of an ASCII string containing the device serial number. • macAddress: Indicates that the value is in the form of Layer 2 MAC address. • other: Indicates that the value is in the form of a platform specific ASCII string containing info that identifies the device. For example: ASCII string contains serialNumber appended/prepended with system name.

ISDP Cache Table

From the **ISDP Cache Table** page, you can view information about other devices the switch has discovered through the ISDP.

To access this page, click **System > Advanced Configuration > ISDP > Cache Table** in the navigation menu.

The following table describes the fields available on this page.

Table 67: ISDP Cache Table Fields

Field	Description
Device ID	Displays the string with Device ID which is reported in the most recent ISDP message.
Interface	Displays the interface that this neighbor is attached to.
IP Address	The (first) network-layer address that is reported in the Address TLV of the most recently received ISDP message.
Version	Displays the Version string for the neighbor.
Hold Time	Displays the ISDP hold time for the neighbor.
Capability	Displays the ISDP Functional Capabilities for the neighbor.
Platform	Displays the ISDP Hardware Platform for the neighbor.
Port ID	Displays the ISDP port ID string for the neighbor.
Protocol Version	Displays the ISDP Protocol Version for the neighbor.
Last Time Changed	Displays when entry was last modified.
Clear (Button)	Clears all entries from the table. The table is repopulated as ISDP messages are received from neighbors.

ISDP Interface Configuration

From the **ISDP Interface Configuration** page, you can configure the ISDP settings for each interface.

Note



If ISDP is enabled on an interface, it must also be enabled globally in order for the interface to transmit ISDP packets. If the ISDP mode on the **ISDP Global Configuration** page is disabled, the interface will not transmit ISDP packets, regardless of the mode configured on the interface. (See [ISDP Global Configuration](#) on page 87.)

To access this page, click **System > Advanced Configuration > ISDP > Interface** in the navigation menu.

The following table describes the fields available on this page.

Table 68: ISDP Interface Configuration Fields

Field	Description
Interface	Select the interface with the ISDP mode status to configure or view.
ISDP Mode	Use this field to enable or disable ISDP on the selected interface.

ISDP Statistics

From the **ISDP Statistics** page, you can view information about the ISDP packets sent and received by the switch.

To access this page, click **System > Advanced Configuration > ISDP > Statistics** in the navigation menu.

The following table describes the fields available on this page.

Table 69: ISDP Statistics Fields

Field	Description
ISDP Packets Received	Displays the number of all ISDP protocol data units (PDUs) received.
ISDP Packets Transmitted	Displays the number of all ISDP PDUs transmitted.
ISDPv1 Packets Received	Displays the number of v1 ISDP PDUs received.
ISDPv1 Packets Transmitted	Displays the number of v1 ISDP PDUs transmitted.
ISDPv2 Packets Received	Displays the number of v2 ISDP PDUs received.
ISDPv2 Packets Transmitted	Displays the number of v2 ISDP PDUs transmitted.
ISDP Bad Header	Displays the number of ISDP PDUs that were received with bad headers.
ISDP Checksum Error	Displays the number of ISDP PDUs that were received with checksum errors.
ISDP Transmission Failure	Displays the number of ISDP PDUs transmission failures.
Invalid Format ISDP Packets Received	Displays the number of ISDP PDUs that were received with an invalid format.
Table Full	Displays the number of times the system tried to add an entry to the ISDP table but was unsuccessful because the table was full.
ISDP IP Address Table Full	Displays the number of times the system tried to add an entry to the ISDP IP Address table but was unsuccessful because the table was full.
Clear (Button)	Resets all statistics to zero.

Configuring Link Dependency

The link dependency feature provides the ability to enable or disable one or more ports based on the link state of one or more different ports. With link dependency enabled on a port, the link state of that port is dependent on the link state of another port. For example, if port A is dependent on port B and the switch detects a link loss on port B, the switch automatically brings down the link on port A. When the link is restored to port B, the switch automatically restores the link to port A.

Link Dependency Group Status

Use the **Link Dependency Group Status** page to configure link dependency groups. Link dependency allows the link status of one interface to be dependent on the link status of another interface. Link state groups define the interface link dependency.

To access this page, click **System > Advanced Configuration > Link Dependency > Group** in the navigation menu.

Use the buttons to perform the following tasks:

- To add a link dependency group, click **Add**. Then, specify a group number, link action, and the interfaces that share a dependency.
- To change the settings for a group, select the checkbox for the group and click **Edit**.
- To delete a link dependency group, select the checkbox for each entry to delete and click **Remove**.
- To view additional information about a group, select the checkbox for the group and click **Details**.

Table 70: Link Dependency Group Status Fields

Field	Description
Group	The unique link dependency group identifier.
Downstream Interfaces	The set of interfaces that depend on other interfaces. In other words, the link state of the downstream interfaces depends on the link state of the upstream interfaces.
Upstream Interfaces	The set of interfaces that determine the link state of the downstream interfaces.
Link Action	<p>The action performed on downstream interfaces when the upstream interfaces are down, which can be one of the following:</p> <ul style="list-style-type: none"> • Up: Downstream interfaces are up when upstream interfaces are down. • Down: Downstream interfaces go down when upstream interfaces are down. <p>Creating a link dependency group with the up link action essentially creates a backup link for the dependent link and alleviates the need to implement <i>STP (Spanning Tree Protocol)</i> to handle the fail-over.</p>
State	<p>The group state, which can be one of the following:</p> <ul style="list-style-type: none"> • Up: Link action is up, and no upstream interfaces have their link up, or link action is down and there are upstream interfaces that have their link up. • Down: Link is down when the above conditions are not true.
Available Interfaces	<p>Available in the Add Group dialog, this field lists the interfaces that can be added to the group. An interface defined as an upstream interface cannot be defined as a downstream interface in the same link state group or in a different group. Similarly, an interface defined as a downstream interface cannot be defined as an upstream interface.</p> <p>To move an interface between the Available Interfaces and Downstream Interfaces or Upstream Interfaces fields, click the interface (or [Ctrl] + click to select multiple interfaces), and then click the appropriate arrow to move the selected interfaces to the desired field.</p>
Link Up	Available in the Group Entry Details dialog, this field lists the upstream and downstream interfaces that currently have their link up.
Link Down	Available in the Group Entry Details dialog, this field lists the upstream and downstream interfaces that currently have their link down.

Configuring Link Local Protocol Filtering

Link Local Protocol Filtering (LLPF) allows the switch to filter out multiple proprietary protocol PDUs, such as Port Aggregation Protocol (PAgP), if the problems occur with proprietary protocols running on

standards-based switches. If certain protocol PDUs cause unexpected results, LLPF can be enabled to prevent those protocol PDUs from being processed by the switch.



Note

The LLPF feature is not supported on all platforms.

LLPF Interface Configuration

Use the **Link Local Protocol Filtering Configuration** page to enable or disable the filtering of various proprietary protocols.

To access this page, click **System > Advanced Configuration > LLPF > Configuration** in the navigation menu.

Use the buttons to perform the following tasks:

- To select the protocols for LLPF to block on an interface, click **Add**. Then, select an interface to configure and select each protocol to block on that interface.
- To change which protocols are blocked on an interface, select the checkbox for the interface and click **Edit**.
- To delete an entry from the list, select the checkbox for each entry to delete and click **Remove**.

The following table describes the fields available on the Link Local Protocol Filtering Configuration page.

Table 71: Link Local Protocol Filtering Configuration

Field	Description
Interface	Identifies the physical or <i>LAG (Link Aggregation Group)</i> interface.
ISDP	When enabled, the select port blocks ISDP PDUs.
VTP	When enabled, the select port blocks VTP PDUs.
DTP	When enabled, the select port blocks DTP PDUs.
UDLD	When enabled, the select port blocks UDLD PDUs.
PAGP	When enabled, the select port blocks PAgP PDUs.
SSTP	When enabled, the select port blocks SSTP PDUs.
All Protocols	All the above mentioned protocols will be dropped in addition to protocols with a Destination MAC of 01:00:0C:CC:CC:CX.

When you configure the blocked protocols on the **Add LLPF Interface** or **Edit LLPF Interface** page, select the checkbox for each protocol to block, or clear the box to allow the protocol on the selected interface. If you select the All Protocols option, all protocols are blocked whether their associated box is checked or unchecked.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Configuring sFlow

sFlow is the standard for monitoring high-speed switched and routed networks. sFlow technology is built into network equipment and gives complete visibility into network activity, enabling effective management and control of network resources.

The sFlow monitoring system consists of an sFlow Agent (embedded in a switch or router or in a standalone probe) and a central sFlow Collector. The sFlow Agent uses sampling technology to capture traffic statistics from the device it is monitoring. sFlow datagrams are used to immediately forward the sampled traffic statistics to an sFlow Collector for analysis.

The sFlow Agent uses two forms of sampling: statistical packet-based sampling of switched or routed Packet Flows, and time-based sampling of counters.

sFlow Agent Summary

Packet Flow Sampling and Counter Sampling are performed by sFlow Instances associated with individual Data Sources within the sFlow Agent. Packet Flow Sampling and Counter Sampling are designed as part of an integrated system. Both types of samples are combined in sFlow datagrams. Packet Flow Sampling will cause a steady, but random, stream of sFlow datagrams to be sent to the sFlow Collector. Counter samples may be taken opportunistically in order to fill these datagrams.

In order to perform Packet Flow Sampling, an sFlow Sampler Instance is configured with a Sampling Rate. The Packet Flow sampling process results in the generation of Packet Flow Records. In order to perform Counter Sampling, the sFlow Poller Instance is configured with a Polling Interval. The Counter Sampling process results in the generation of Counter Records. The sFlow Agent collects Counter Records and Packet Flow Records and sends them in the form of sFlow datagrams to sFlow Collectors.

To access the **sFlow Agent Summary** page, click **System > Advanced Configuration > sFlow > Agent** in the navigation menu.

Table 72: sFlow Agent Summary Fields

Field	Description
Version	Uniquely identifies the version and implementation of this MIB. The version string must have the following structure: MIB Version; Organization; Software Revision where: <ul style="list-style-type: none"> • MIB Version: '1.3', the version of this MIB. • Organization: Broadcom Corp. • Revision: 1.0
Agent Address	The IP address associated with this agent.

Use the **Refresh** button to refresh the page with the most current data from the switch.

sFlow Receiver Configuration

Use the **sFlow Receiver Configuration** page to configure the sFlow Receiver.

To access this page, click **System > Advanced Configuration > sFlow > Receiver** in the navigation menu.

Table 73: sFlow Receiver Configuration Fields

Field	Description
Index	Selects the receiver for which data is to be displayed or configured. The allowed range is 1 to 8.
Owner String	The entity making use of this sFlowRcvrTable entry. The empty string indicates that the entry is currently unclaimed and the receiver configuration is reset to the default values. An entity wishing to claim an sFlowRcvrTable entry must ensure that the entry is unclaimed before trying to claim it. The entry is claimed by setting the owner string. The entry must be claimed before any changes can be made to other sampler objects.
Time Remaining	The time (in seconds) remaining before the sampler is released and stops sampling. A value of 0 essentially means the receiver is not configured. When configuring the sFlow receiver settings, you must select the Timeout Mode option before you can configure a Timeout Value.
Maximum Datagram Size	The maximum number of data bytes that can be sent in a single sample datagram. The manager should set this value to avoid fragmentation of the sFlow datagrams. The default value is 1400. The allowed range is 200 to 9116.
Address	The IP address of the sFlow collector. If set to 0.0.0.0 no sFlow datagrams will be sent.
Port	The destination port for sFlow datagrams. The allowed range is 1 to 65535.
Datagram Version	The version of sFlow datagrams that should be sent.
Monitor Session	Monitor session to enable sFlow hardware feature.

Use the **Submit** button to sent updated data to the switch and cause the changes to take effect on the switch.

Use the **Refresh** button to refresh the page with the most current data from the switch.

Use the **Edit** button to configure the monitor session for a specific receiver (only for IPv4). After successful configuration, the sFlow packet processing will be done in hardware.

Use the **Submit** button to sent updated data to the switch and cause the changes to take effect on the switch.

sFlow Poller Configuration

The sFlow agent collects time-based sampling of network interface statistics and sends them to the configured sFlow receivers. A data source configured to collect counter samples is called a poller.

Counter Sampling

Use the **sFlow Poller Configuration** page to efficiently, periodically export counters associated with Data Sources. A maximum Sampling Interval is assigned to each sFlow instance associated with a Data Source.

Counter Sampling is accomplished as follows:

The sFlow Agent keeps a list of counter sources being sampled. When a Packet Flow Sample is generated, the sFlow Agent examines the list and adds counters to the sample datagram, least recently sampled first. Counters are added to the datagram only when the sources are within five seconds of

failing to meet the required Sampling Interval. Once a second, the sFlow Agent examines the list of counter sources and sends any counters that need to be sent to meet the sampling interval requirement.

To access this page, click **System > Advanced Configuration > sFlow > Poller** in the navigation menu.

Use the buttons to perform the following tasks:

- To add an sFlow poller instance, click **Add** and complete the required information.
- To edit an existing sFlow poller instance, select the appropriate checkbox or click the row to select the sFlow poller instance and click **Edit**. Modify the sFlow poller configuration information as needed.
- To delete an sFlow poller instance, select one or more table entries and click **Remove**.

Table 74: sFlow Poller Configuration Fields

Field	Description
Poller DataSource	The sFlow Sampler Datasource for this flow sampler. This Agent will support physical ports only.
Receiver Index	The sFlowReceiver for this sFlow Counter Poller. If set to zero, the poller configuration is set to the default and the poller is deleted. Only active receivers can be set. If a receiver expires, then all pollers associated with the receiver will also expire. The allowed range is 1 to 8.
Poller Interval	The maximum number of seconds between successive samples of the counters associated with this data source.

Click **Refresh** to refresh the page with the most current data from the switch.

sFlow Sampler Configuration

The sFlow Agent collects a statistical packet-based sampling of the switched flows and sends them to the configured receivers. A data source configured to collect flow samples is called a sampler.

Packet Flow Sampling

The Packet Flow Sampling mechanism carried out by each sFlow instance ensures that any packet observed at a Data Source has an equal chance of being sampled, irrespective of the Packet Flow(s) to which it belongs.

Packet Flow Sampling is accomplished as follows:

- When a packet arrives on an interface, the Network Device makes a filtering decision to determine whether the packet should be dropped.
- If the packet is not filtered (dropped), a destination interface is assigned by the switching/routing function.
- At this point, a decision is made on whether or not to sample the packet. The mechanism involves a counter that is decremented with each packet. When the counter reaches zero, a sample is taken. When a sample is taken, the counter that indicates how many packets to skip before taking the next sample is reset. The value of the counter is set to a random integer where the sequence of random integers used over time is the Sampling Rate.

To access the **sFlow Sampler Configuration** page, click **System > Advanced Configuration > sFlow > Sampler** in the navigation menu.

Table 75: sFlow Sampler Configuration Fields

Field	Description
Sampler Data Source	The sFlowDataSource for this sFlow sampler. The sFlow agent supports physical ports as sFlow data sources.
Receiver Index	The sFlowReceiver for this sFlow sampler. The specified Receiver Index must be associated with an active sFlow receiver. If a receiver expires, all samplers associated with the receiver will also expire.
Remote Agent Index	The remote agent index.
Ingress Sampling Rate	sFlow instance packet Sampling Rate for Ingress sampling.
Egress Sampling Rate	sFlow instance egress packet Sampling Rate.
Flow-based Sampling Rate	sFlow instance flow based packet Sampling Rate.
Maximum Header Size	The maximum number of bytes that should be copied from a sampled packet.
IP ACL	The ID of the IP ACL to apply to traffic from the sampler.
IP MAC	The ID of the MAC ACL to apply to traffic from the sampler.

Click **Refresh** to refresh the page with the most current data from the switch.

Use the buttons to perform the following tasks:

- To add an sFlow sampler instance, click **Add** and complete the required information.
- To edit an existing sFlow sampler instance, select the appropriate checkbox or click the row to select the sFlow sampler instance and click **Edit**. Modify the sFlow sampler configuration information as needed.
- To delete an sFlow sampler instance, select one or more table entries and click **Remove**.

The **Add Sampler** page lets you configure the sampling rate for ingress/egress/flow based sampling. After successful configuration, the sFlow packet sampling is performed based on sampling rate.

Table 76: sFlow Sampler Configuration Fields

Field	Description
Sampler Data Source	The sFlowDataSource for this sFlow sampler. The sFlow agent supports physical ports as sFlow data sources.
Receiver Index	The sFlow Receiver for this sFlow sampler. If set to zero, no packets will be sampled. Only active receivers can be set. If a receiver expires, then all samplers associated with the receiver will also expire. The allowed range is 1 to 8.
Ingress Sampling Rate	sFlow instance packet Sampling Rate for Ingress sampling.

Table 76: sFlow Sampler Configuration Fields (continued)

Field	Description
Egress Sampling Rate	sFlow instance egress packet Sampling Rate.
Flow-based Sampling Rate	sFlow instance flow based packet Sampling Rate.
Maximum Header Size	The maximum number of bytes that should be copied from a sampled packet. The allowed range is 20 to 256.

sFlow Source Interface Configuration

Use the **sFlow Source Interface Configuration** page to specify the physical or logical interface to use as the sFlow client source interface. When an IP address is configured on the source interface, this address is used for all sFlow communications between the local sFlow client and the remote sFlow server. The IP address of the designated source interface is used in the IP header of sFlow management protocol packets. This allows security devices, such as firewalls, to identify all source packets coming from a specific device.

To access this page, click **System > Advanced Configuration > sFlow > Source Interface Configuration** in the navigation menu.

Table 77: sFlow Source Interface Configuration Fields

Field	Description
Type	The type of interface to use as the source interface: <ul style="list-style-type: none"> • None: The primary IP address of the originating (outbound) interface is used as the source address. • Interface: The primary IP address of a physical port is used as the source address. • Loopback: The primary IP address of the loopback interface is used as the source address. A loopback is always reachable, as long as any routing interface is up. • VLAN - The primary IP address of a VLAN routing interface is used as the source address.
Interface	When the selected Type is Interface, select the physical port to use as the source interface.
Loopback Interface	When the selected Type is Loopback, select the loopback interface to use as the source interface.
VLAN ID	When the selected Type is VLAN, select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.
Tunnel ID	The primary IP address of a tunnel interface is used as the source address.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Configuring SNTP Settings

200 Series software supports the *SNTP*. SNTP assures accurate network device clock time synchronization up to the millisecond. Time synchronization is performed by a network SNTP server. 200 Series software operates only as an SNTP client and cannot provide time services to other systems.

Time sources are established by Stratum. Stratum defines the accuracy of the reference clock. The higher the stratum (where zero is the highest), the more accurate the clock. The device receives time from stratum 1 and above since it is itself a stratum 2 device.

The following is an example of stratum:

- **Stratum 0:** A real time clock is used as the time source, for example, a GPS system.
- **Stratum 1:** A server that is directly linked to a Stratum 0 time source is used. Stratum 1 time servers provide primary network time standards.
- **Stratum 2:** The time source is distanced from the Stratum 1 server over a network path. For example, a Stratum 2 server receives the time over a network link, via NTP, from a Stratum 1 server.

Information received from SNTP servers is evaluated based on the time level and server type.

SNTP time definitions are assessed and determined by the following time levels:

- **T1:** Time at which the original request was sent by the client.
- **T2:** Time at which the original request was received by the server.
- **T3:** Time at which the server sent a reply.
- **T4:** Time at which the client received the server's reply.

The device can poll unicast and broadcast server types for the server time.

Polling for unicast information is used for polling a server for which the IP address is known. SNTP servers that have been configured on the device are the only ones that are polled for synchronization information. T1 through T4 are used to determine server time. This is the preferred method for synchronizing device time because it is the most secure method. If this method is selected, SNTP information is accepted only from SNTP servers defined on the device using the SNTP Server Configuration page.

Broadcast information is used when the server IP address is unknown. When a broadcast message is sent from an SNTP server, the SNTP client listens to the message. If broadcast polling is enabled, any synchronization information is accepted, even if it has not been requested by the device. This is the least secure method.

The device retrieves synchronization information, either by actively requesting information or at every poll interval. If unicast and broadcast polling are enabled, the information is retrieved in this order:

- Information from servers defined on the device is preferred. If unicast polling is not enabled or if no servers are defined on the device, the device accepts time information from any SNTP server that responds.
- If more than one unicast device responds, synchronization information is preferred from the device with the lowest stratum.
- If the servers have the same stratum, synchronization information is accepted from the SNTP server that responded first.

MD5 (Message-Digest algorithm 5) Authentication safeguards device synchronization paths to SNTP servers. MD5 is an algorithm that produces a 128-bit hash. MD5 is a variation of MD4, and increases MD4 security. MD5 verifies the integrity of the communication, authenticates the origin of the communication.

SNTP Global Configuration

Use the **SNTP Global Configuration** page to view and adjust *SNTP* parameters.

To display this page, click **System > Advanced Configuration > SNTP > Global Configuration** in the navigation menu.

Table 78: SNTP Global Configuration Fields

Field	Description
Client Mode	Use the drop-down menu to specify the SNTP client mode, which is one of the following modes: <ul style="list-style-type: none"> • Disable: SNTP is not operational. No SNTP requests are sent from the client nor are any received SNTP messages processed. • Unicast: SNTP operates in a point to point fashion. A unicast client sends a request to a designated server at its unicast address and expects a reply from which it can determine the time and, optionally the round-trip delay and local clock offset relative to the server. • Broadcast: SNTP operates in the same manner as multicast mode but uses a local broadcast address instead of a multicast address. The broadcast address has a single subnet scope while a multicast address has Internet wide scope.
Port	Specifies the local UDP port to listen for responses/broadcasts. The allowed range is 1 to 65535. The default value is 123.
Unicast Poll Interval	Specifies the number of seconds between unicast poll requests expressed as a power of two when configured in unicast mode. The allowed range is 6 to 10. The default value is 6.
Broadcast Poll Interval	Specifies the number of seconds between broadcast poll requests expressed as a power of two when configured in broadcast mode. Broadcasts received prior to the expiry of this interval are discarded. The allowed range is 6 to 10. The default value is 6.
Unicast Poll Timeout	Specifies the number of seconds to wait for an SNTP response when configured in unicast mode. The allowed range is 1 to 30. The default value is 5.
Unicast Poll Retry	Specifies the number of times to retry a request to an SNTP server after the first time-out before attempting to use the next configured server when configured in unicast mode. The allowed range is 0 to 10. The default value is 1.
Number of Servers Configured	Specifies the number of current valid unicast server entries configured for this client.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

SNTP Global Status

Use the **SNTP Global Status** page to view information about the system's *SNTP* client.

To access this page, click **System > Advanced Configuration > SNTP > Global Status** in the navigation menu.

Table 79: Global Status Fields

Field	Description
Version	Specifies the SNTP Version the client supports.
Supported Mode	Specifies the SNTP modes the client supports. Multiple modes may be supported by a client.
Last Update Time	Specifies the local date and time (UTC) the SNTP client last updated the system clock.
Last Attempt Time	Specifies the local date and time (UTC) of the last SNTP request or receipt of an unsolicited message.
Last Attempt Status	<p>Specifies the status of the last SNTP request or unsolicited message for both unicast and broadcast modes. If no message has been received from a server, a status of Other is displayed. These values are appropriate for all operational modes:</p> <ul style="list-style-type: none"> • Other: None of the following enumeration values. • Success: The SNTP operation was successful and the system time was updated. • Request Timed Out: A directed SNTP request timed out without receiving a response from the SNTP server. • Bad Date Encoded: The time provided by the SNTP server is not valid. • Version Not Supported: The SNTP version supported by the server is not compatible with the version supported by the client. • Server Unsynchronized: The SNTP server is not synchronized with its peers. This is indicated via the 'leap indicator' field on the SNTP message. • Server Kiss Of Death: The SNTP server indicated that no further queries were to be sent to this server. This is indicated by a stratum field equal to 0 in a message received from a server.
Server IP Address	Specifies the IP address of the server for the last received valid packet. If no message has been received from any server, an empty string is shown.
Address Type	Specifies the address type of the SNTP Server address for the last received valid packet.
Server Stratum	Specifies the claimed stratum of the server for the last received valid packet.
Reference Clock Id	Specifies the reference clock identifier of the server for the last received valid packet.
Server Mode	Specifies the mode of the server for the last received valid packet.
Unicast Sever Max Entries	Specifies the maximum number of unicast server entries that can be configured on this client.
Unicast Server Current Entries	Specifies the number of current valid unicast server entries configured for this client.
Broadcast Count	Specifies the number of unsolicited broadcast SNTP messages that have been received and processed by the SNTP client since last reboot.

Click **Refresh** to display the latest information from the router.

SNTP Server Configuration

Use the **SNTP Server Configuration** page to view and modify information for adding and modifying *SNTP* servers.

To display this page, click **System > Advanced Configuration > SNTP > Server Configuration** in the navigation menu.

Table 80: SNTP Server Configuration Fields

Field	Description
SNTP Server	Select the IP address of a user-defined SNTP server to view or modify information about an SNTP server, or select Add to configure a new SNTP server. You can define up to three SNTP servers.
Type	Select IPv4 if you entered an IPv4 address, IPv6 if you entered an IPv6 address or DNS if you entered a hostname.
Port	Enter a port number from 1 to 65535. The default is 123.
Priority	Enter a priority from 1 to 3, with 1 being the highest priority. The switch will attempt to use the highest priority server and, if it is not available, will use the next highest server.
Version	Enter the protocol version number.

- To add an SNTP server, select **Add** from the **Server** list, complete the remaining fields as desired, and click **Submit**. The SNTP server is added, and is now reflected in the Server list. You must perform a save to retain your changes over a power cycle.
- To remove an SNTP server, select the IP address of the server to remove from the **Server** list, and then click **Remove**. The entry is removed, and the device is updated.

SNTP Server Status

The **SNTP Server Status** page displays status information about the *SNTP* servers configured on your switch.

To access this page, click **System > Advanced Configuration > SNTP > Server Status** in the navigation menu.

Table 81: SNTP Server Status Fields

Field	Description
Address	Specifies all the existing Server Addresses. If no Server configuration exists, <code>No SNTP server exists</code> displays.
Last Update Time	Specifies the local date and time (UTC) that the response from this server was used to update the system clock.
Last Attempt Time	Specifies the local date and time (UTC) that this SNTP server was last queried.

Table 81: SNTP Server Status Fields (continued)

Field	Description
Last Attempt Status	Specifies the status of the last SNTP request to this server. If no packet has been received from this server, a status of Other is displayed: <ul style="list-style-type: none"> • Other: None of the following enumeration values. • Success: The SNTP operation was successful and the system time was updated. • Request Timed Out: A directed SNTP request timed out without receiving a response from the SNTP server. • Bad Date Encoded: The time provided by the SNTP server is not valid. • Version Not Supported: The SNTP version supported by the server is not compatible with the version supported by the client. • Server Unsynchronized: The SNTP server is not synchronized with its peers. This is indicated via the 'leap indicator' field on the SNTP message. • Server Kiss Of Death: The SNTP server indicated that no further queries were to be sent to this server. This is indicated by a stratum field equal to 0 in a message received from a server.
Requests	Specifies the number of SNTP requests made to this server since last agent reboot.
Failed Requests	Specifies the number of failed SNTP requests made to this server since last reboot.

Click **Refresh** to display the latest information from the switch.

SNTP Source Interface Configuration

Use the **SNTP Source Interface Configuration** page to specify the physical or logical interface to use as the *SNTP* client source interface. When an IP address is configured on the source interface, this address is used for all SNTP communications between the local SNTP client and the remote SNTP server. The IP address of the designated source interface is used in the IP header of SNTP management protocol packets. This allows security devices, such as firewalls, to identify all source packets coming from a specific device.

To access this page, click **System > Advanced Configuration > SNTP > Source Interface Configuration** in the navigation menu.

Table 82: SNTP Source Interface Configuration Fields

Field	Description
Type	The type of interface to use as the source interface: <ul style="list-style-type: none"> • None: The primary IP address of the originating (outbound) interface is used as the source address. • Interface: The primary IP address of a physical port is used as the source address. • Loopback: The primary IP address of the loopback interface is used as the source address. A loopback is always reachable, as long as any routing interface is up. • VLAN: The primary IP address of a VLAN routing interface is used as the source address.
Interface	When the selected Type is Interface, select the physical port to use as the source interface.

Table 82: SNTP Source Interface Configuration Fields (continued)

Field	Description
Loopback Interface	When the selected Type is Loopback, select the loopback interface to use as the source interface.
VLAN ID	When the selected Type is VLAN, select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.
Tunnel ID	When the selected Type is Tunnel, select the tunnel interface to use as the source interface.

Click **Refresh** to display the latest information from the switch.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Configuring Time Ranges

You can use these pages to configure time ranges to use in time-based *ACL (Access Control List)* rules. Time-based ACLs allow one or more rules within an ACL to be based on a periodic or absolute time. Each ACL rule within an ACL except for the implicit deny all rule can be configured to be active and operational only during a specific time period. The time range pages allow you to define specific times of the day and week in order to implement time-based ACLs. The time range is identified by a name and can then be referenced by an ACL rule defined within an ACL.

Time Range Configuration

Use the **Time Range Configuration** page to create a named time range. Each time range can consist of one absolute time entry and/or one or more periodic time entries.

To access this page, click **System > Advanced Configuration > Time Ranges > Configuration** in the navigation menu.

Table 83: Time Range Configuration

Field	Description
Time Range Name	The unique ID or name that identifies this time range. A time-based <i>ACL</i> rule can reference the name configured in this field.
Time Range Status	Shows whether the time range is Active or Inactive. A time range is Inactive if the current day and time do not fall within any time range entries configured for the time range.
Periodic Entry Count	The number of periodic time range entries currently configured for the time range.
Absolute Entry	Shows whether an absolute time entry is currently configured for the time range.

Use the buttons to perform the following tasks:

- To add a time range, click **Add** and configure a name for the time range configuration.
- To delete a configured time range, select each entry to delete, click **Remove**, and confirm the action.
- Use **Submit** to add a new time range.

Time Range Entry Configuration

Use the **Time Range Entry Configuration** page to configure periodic and absolute time range entries and add them to named time ranges.



Note

The time range entries use the system time for the time periods in which they take effect. Make sure you configure the SNTP server settings so that the SNTP client on the switch can obtain the correct date and time from the server.

To access this page, click **System > Advanced Configuration > Time Ranges > Entry Configuration** in the navigation menu.

To configure the time range entries for a time range configuration, select the time range configuration from the Time Range Name menu and use the buttons to perform the following tasks:

- To add an Absolute time range entry, click **Add Absolute** and configure information about when the Absolute entry occurs. If the **Add Absolute** button is not available, an Absolute entry already exists for the selected time range configuration.
- To add a Periodic time range entry, click **Add Periodic** and specify the days and times that the entry is in effect.
- To delete a time range entry, select each entry to delete, click **Remove**, and confirm the action.

Table 84: Time Range Entry Configuration

Field	Description
Time Range Name	Select the name of the time range to which you want to add a time range entry.
Time Range Entry	Select Create New Time Range Entry to add a new entry to a time range. To view or delete an existing time range entry, select its ID from the menu.
Time Range Entry ID	When creating a new time range entry, assign a unique ID number from 1-10. This field does not appear if the entry has already been configured.
Time Range Entry Type.	Specifies whether the entry is periodic or absolute. A periodic entry occurs at the same time every day or on one or more days of the week. An absolute entry does not repeat.
Periodic Time Range Entry	
Applicable Days	Specify the day(s) when the time entry occurs: <ul style="list-style-type: none"> • Daily: Has the same start and end time every day • Weekdays: Has the same start and end time Monday through Friday • Weekends: Has the same start and end time on Saturday and Sunday • Days of the Week: Select the day of the week when the entry starts and stops. You do not need to use the same day of the week for the start and end time.
Start Day	(Periodic Days of Week only) Select the day the time range entry starts. To select multiple days, hold the [Ctrl] key and click the days.
Start Time	Specify the time when the entry begins. The time is based on a 24-hour clock. For example, 6:00 PM is 18:00.
End Day	(Periodic Days of Week only) Select the day the time range entry ends.

Table 84: Time Range Entry Configuration (continued)

Field	Description
End Time	Specify the time when the entry ends. The time is based on a 24-hour clock. For example, 6:00 PM is 18:00.
Absolute Time Range Entry	
Absolute Start Date and Time	Select the checkbox to configure the date and time when the time range entry begins.
Start Month	Select the month when the time entry begins.
Start Date	Select the day of the month when the time entry begins.
Start Year	Select the year when the time entry begins.
Start Time	Specify the time when the entry begins. The time is based on a 24-hour clock. For example, 6:00 PM is 18:00.
Absolute End Date and Time	Select the checkbox to configure the date and time when the time range entry ends.
End Month	Select the month when the time entry ends.
End Date	Select the day of the month when the time entry ends.
End Year	Select the year when the time entry ends.
End Time	Specify the time when the entry ends. The time is based on a 24-hour clock. For example, 6:00 PM is 18:00.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Configuration changes take effect immediately.

Configuring the Time Zone

The **Time Zone Summary** page displays information about the current system time, the time zone, and the daylight saving time (also known as summer time) settings configured on the device.

To access this page, click **System > Advanced Configuration > Time Zone > Summary** in the navigation menu.

Table 85: Time Zone Summary Fields

Field	Description
Current Time	<p>This section contains information about the system time and date on the device. If the current time has not been acquired by the <i>SNTP</i> client on the device or configured manually, this section shows the default time and date plus the amount of time since the system was last rebooted.</p> <ul style="list-style-type: none"> • Time: The current time on the system clock. This time is used to provide time stamps on log messages. Additionally, some CLI show commands include the time in the command output. • Zone: The acronym that represents the time zone. • Date: The current date on the system. • Time Source: The time source from which the time update is taken – one of the following: <ul style="list-style-type: none"> • SNTP: The time has been acquired from an SNTP server. • No Time Source: The time has either been manually configured or not configured at all.
Time Zone	<p>This section contains information about the time zone and offset.</p> <ul style="list-style-type: none"> • Zone: The acronym that represents the time zone. It can be any 0- to 4-character text string. • Offset: The number of hours offset from Coordinated Universal Time (UTC), which is also known as Greenwich Mean Time (GMT).

Click **Refresh** to display the latest information from the router.

Time Zone Configuration

Use the **Time Zone Configuration** page to manually configure the system clock settings. The *SNTP* client must be disabled to allow manual configuration of the system time and date.

To access this page, click **System > Advanced Configuration > Time Zone > Time Zone** in the navigation menu.

Table 86: Time Zone Configuration Fields

Field	Description
Time Zone	<p>The time zone settings include the amount of time the system clock is offset from Coordinated Universal Time (UTC) and the time zone acronym.</p> <ul style="list-style-type: none"> • Offset: The number of hours the system clock is offset from UTC, which is also known as Greenwich Mean Time (GMT). • Zone: The acronym that represents the time zone. It can be any 0- to 4-character text string.
Date and Time	<p>Use the fields in this section to manually configure the system time and date. If the SNTP client is enabled (unicast or broadcast mode), these fields cannot be configured.</p> <ul style="list-style-type: none"> • Time: The current time in hours, minutes, and seconds on the system clock. • Date: The current date in month, day, and year on the system clock. To change the date, click the calendar icon to the right of the field, select the year from the menu, browse to the desired month, and click the date.

Click **Refresh** to display the latest information from the router.

Click **Submit** to apply the settings to the running configuration and cause the change to take effect.

Summer Time Configuration

Use the **Summer Time Configuration** page to configure settings for summer time, which is also known as daylight saving time. Used in some countries around the world, summer time is the practice of temporarily advancing clocks during the summer months. Typically clocks are adjusted forward one or more hours near the start of spring and are adjusted backward in autumn.

To access this page, click **System > Advanced Configuration > Time Zone > Summer Time** in the navigation menu.

Table 87: Summer Time Configuration Fields

Field	Description
Summer Time	<p>The summer time mode on the system:</p> <ul style="list-style-type: none"> • Disable: Summer time is not active, and the time does not shift based on the time of year. • Recurring: Summer time occurs at the same time every year. The start and end times and dates for the time shift must be manually configured. • EU: The system clock uses the standard recurring summer time settings used in countries in the European Union. When this field is selected, the rest of the applicable fields on the page are automatically populated and cannot be edited. • USA: The system clock uses the standard recurring daylight saving time settings used in the United States. When this field is selected, the rest of the applicable fields on the page are automatically populated and cannot be edited. • Non-Recurring: Summer time settings are in effect only between the start date and end date of the specified year. When this mode is selected, the summer time settings do not repeat on an annual basis.
Date Range	<p>The fields in this section are available only if the Non-Recurring mode is selected from the Summer Time menu.</p> <ul style="list-style-type: none"> • Start Date: The day, month, and year that summer time begins. To change the date, click the calendar icon to the right of the field, select the year from the menu, browse to the desired month, and click the date. • Starting Time of Day: The time, in hours and minutes, to start summer time on the specified day. • End Date: The day, month, and year that summer time ends. To change the date, click the calendar icon to the right of the field, select the year from the menu, browse to the desired month, and click the date. • Ending Time of Day: The time, in hours and minutes to end summer time on the specified day.

Table 87: Summer Time Configuration Fields (continued)

Field	Description
Recurring Date	<p>The fields in this section are available only if the Recurring mode is selected from the Summer Time menu.</p> <ul style="list-style-type: none"> • Start Week: The week of the month within which summer time begins. • Start Day: The day of the week on which summer time begins. • Start Month: The month of the year within which summer time begins. • Starting Time of Day: The time, in hours and minutes, to start summer time. • End Week: The week of the month within which summer time ends. • End Day: The day of the week on which summer time ends. • End Month: The month of the year within which summer time ends. • Ending Time of Day: The time, in hours and minutes, to end summer time.
Zone	<p>The fields in this section are available only if the Recurring or Non-Recurring modes are selected from the Summer Time menu.</p> <ul style="list-style-type: none"> • Offset: The number of minutes to shift the summer time from the standard time. • Zone: The acronym associated with the time zone when summer time is in effect.

Click **Refresh** to display the latest information from the router.

Click **Submit** to apply the settings to the running configuration and cause the change to take effect.

Managing SNMP Traps

The pages in the Trap Manager folder allow you to view and configure information about [SNMP](#) traps the system generates.

Trap Log

Use the **Trap Log** page to view the entries in the trap log.

To access this page, click **System > Advanced Configuration > Trap Manager > Trap Log** in the navigation menu.

Table 88: Trap Log Fields

Field	Description
Trap Log Capacity	The maximum number of traps stored in the log. If the number of traps exceeds the capacity, the entries will overwrite the oldest entries.
Number of Traps Since Last Reset	The number of traps generated since the trap log entries were last cleared.
Number of Traps Since Log Last Viewed	The number of traps that have occurred since the traps were last displayed. Displaying the traps by any method (terminal interface display, web display, upload file from switch, etc.) will cause this counter to be cleared to 0.
Log	The sequence number of this trap.

Table 88: Trap Log Fields (continued)

Field	Description
System Up Time	The time at which this trap occurred, expressed in days, hours, minutes and seconds since the last reboot of the switch.
Trap	Displays the information identifying the trap.

Click **Clear Log** to clear all entries in the log. Subsequent displays of the log will only show new log entries.

Trap Flags

Use the **Trap Flags** page to enable or disable traps the switch can send to an *SNMP* manager. When the condition identified by an active trap is encountered by the switch, a trap message is sent to any enabled SNMP Trap Receivers, and a message is written to the trap log.

To access this page, click **System > Advanced Configuration > Trap Manager > Trap Flags** in the navigation menu.

The fields available on the Trap Flags page depend on the packages installed on your system. For example, if your system does not have the BGP4 package installed, the **BGP Traps** field is not available. [Table 89](#) shows the fields that are available on a system with all packages installed.

Table 89: Trap Flags Configuration Fields

Field	Description
Authentication	Enable or disable activation of authentication failure traps by selecting the corresponding line on the drop-down entry field. The factory default is enabled.
Link Up/Down	Enable or disable activation of link status traps by selecting the corresponding line on the drop-down entry field. The factory default is enabled.
Multiple Users	Enable or disable activation of multiple user traps by selecting the corresponding line in the drop-down menu. The factory default is enabled. This trap is triggered when the same user ID is logged into the switch more than once at the same time (either via Telnet or the serial port).
Spanning Tree	Enable or disable activation of spanning tree traps by selecting the corresponding line on the drop-down entry field. The factory default is enabled.
ACL Traps	Enable or disable activation of <i>ACL</i> traps by selecting the corresponding line on the drop-down entry field. The factory default is disabled.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Managing CPU Traffic Filters

The pages in the CPU Traffic Filter folder allow you to configure CPU traffic filtering and view data about filtered traffic.

CPU Traffic Filter Global Configuration

Use the **CPU Traffic Filter Global Configuration** page to view and modify the CPU Traffic Filter settings on the device.

To access this page, click **System > Advanced Configuration > CPU Traffic Filter > Global** in the navigation menu.

Table 90: CPU Traffic Filter Global Configuration Fields

Field	Description
Admin Mode	Enables CPU-traffic mode. The packets in the Tx (transmitted) and Rx (received) directions are matched when the mode is enabled. The default value is disabled.
CPU Trace Mode	Enables CPU packet tracing. The packet may be received by multiple components. If CPU packet tracing is enabled and tracing is configured, then the packets are traced according to the defined filter.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Click **Refresh** to refresh the page with the most current data from the switch.

CPU Traffic Filter Configuration

Use the **CPU Traffic Filter Configuration** page to create, edit, or remove CPU traffic filters and to view summary information about the filters that exist on the device.

To access this page, click **System > Advanced Configuration > CPU Traffic Filter > Filter Configuration** in the navigation menu.

Use the buttons to perform the following tasks:

- To edit an existing filter for a direction, select the entry to modify and click **Edit**.
- To edit CPU traffic filters for both directions, select the **Tx** (transmitted) and **Rx** (received) checkboxes and click **Edit**.
- To remove one or more configured filters, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.

Table 91: CPU Traffic Filter Configuration Fields

Field	Description
Direction	Two software filters are used (one filter in each direction – Tx, Rx, or both) with condition matching as one, many or all in the below list in the specified direction.
Filter	<p>Specific filters for each direction. The filter options are:</p> <p>Protocol: Specific protocol filters. The statistics and/or traces for configured filters are obtained for the packet matching the configured filter. The protocol options are:</p> <ul style="list-style-type: none"> • STP • LACPDU • ARP • UDLD • LLDP • IP • OSPF • BGP • DHCP • BCAST • MCAST • UCAST • Source IP • Destination IP • Source MAC • Destination MAC • Custom • Source TCP • Destination TCP • Source UDP • Destination UDP <p>IP Address: Source or Destination IP address specific filter. The statistics and/or the traces for configured filters are obtained for the packet matching the configured Source / Destination IP/Mask.</p> <p>MAC Address: Source / Destination MAC address specific filter. The statistics and/or the traces for configured filters are obtained for the packet matching configured Source / Destination MAC address.</p> <p>Custom: Custom filter. The statistics and/or the traces for configured filters are obtained for the packet matching configured data at specific offset. If the mask is not specified, the default mask is 0xFF.</p> <p>TCP Port: Source / Destination TCP Port specific filter. The statistics and/or the traces for configured filters are obtained for the packet matching configured Source / Destination TCP Port.</p> <p>UDP Port: Source / Destination UDP Port specific filter. The statistics and/or the traces for configured filters are obtained for the packet matching configured Source / Destination UDP Port.</p>

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

CPU Traffic Filter Interface Configuration

Use the **CPU Traffic Filter Interface Configuration** page to associate the CPU filters to an interface or a list of interfaces. Each interfaces can be a physical or logical *LAG*. The statistics counters are updated only for the configured interfaces. Similarly, the traces can also be obtained for configured interfaces.

To access this page, click **System > Advanced Configuration > CPU Traffic Filter > Interfaces** in the navigation menu.

Use the buttons to perform the following tasks:

- To add CPU Traffic filter to interface(s), click **Add**.
- To remove one or more associated filters, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.
-

Table 92: CPU Traffic Filter Interface Configuration Fields

Field	Description
Interface	Each interfaces can be a physical or logical <i>LAG</i> .
Direction	Two software filters are used (one filter in each direction – Tx, Rx, or both) with condition matching as one, many or all.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Click **Refresh** to refresh the page with the most current data from the switch.

CPU Traffic Filter Statistics

Use the **CPU Traffic Filter Statistics** page to view per interface statistics for configured CPU filters.

To access this page, click **System > Advanced Configuration > CPU Traffic Filter > Statistics** in the navigation menu.

To clear interface statistics, click **Clear**. You must confirm the action before the data is cleared.

Table 93: CPU Traffic Filter Statistics Fields

Field	Description
Filter Name	The list of available filter names. Select a filter to view its interface statistics.
Interface	Each interface can be a physical or logical <i>LAG</i> .
Tx	The counter statistics for an interface associated with the Tx (transmitted) direction.
Last Updated Tx Timestamp	The time when the sent packet count on a Tx interface was last updated, based on the user-defined packet filter on the interface.

Table 93: CPU Traffic Filter Statistics Fields (continued)

Field	Description
Rx	The counter statistics for an interface associated with the Rx (received) direction.
Last Updated Rx Timestamp	The time when the sent packet count on an Rx interface was last updated, based on the user-defined packet filter on the interface.

Click **Refresh** to refresh the page with the most current data from the switch.

CPU Traffic Filter Summary

Use the **CPU Traffic Filter Summary** page to view a summary of all interfaces for CPU filters.

To access this page, click **System > Advanced Configuration > CPU Traffic Filter > Summary** in the navigation menu.

To clear the filter summary, click **Clear**. You must confirm the action before the data is cleared.

Table 94: CPU Traffic Filter Summary Fields

Field	Description
Filter Name	The associated filter name.
Transmitted	The counter statistics for all interfaces associated with the Tx direction.
Received	The counter statistics for all interfaces associated with the Rx direction.

Click **Refresh** to refresh the page with the most current data from the switch.

CPU Traffic Filter Trace Information

Use the **CPU Traffic Filter Trace Information** page to view CPU trace information.

To access this page, click **System > Advanced Configuration > CPU Traffic Filter > Trace Information** in the navigation menu.

To clear the trace information, click **Clear**. You must confirm the action before the data is cleared.

Table 95: CPU Traffic Filter Trace Information Fields

Field	Description
Trace Information	Trace information for the matching packets as defined in the filters until the packet is delivered to registered application. .

Click **Refresh** to refresh the page with the most current data from the switch.

Viewing the System Firmware Status

The pages in the Firmware folder allow you to view and monitor the system firmware status. The Firmware folder has links to the following pages.

Dual Image Status

The **Dual Image Status** feature allows the switch to have two 200 Series software images in the permanent storage. One image is the active image, and the second image is the backup. This feature reduces the system down-time during upgrades and downgrades. You can use the Dual Image Status page to view information about the system images on the device.

To access this page, click **System > Firmware > Status** in the navigation menu.

Table 96: Dual Image Status Fields

Field	Description
Unit	Displays the unit ID of the switch.
Active	Displays the version of the active code file.
Backup	Displays the version of the backup code file.
Current Active	Displays the currently active image on this unit.
Next Active	Displays the image to be used on the next restart of this unit.
Active Description	Displays the description associated with the active code file.
Backup Description	Displays the description associated with the backup code file.

Click **Refresh** to display the latest information from the router.

For information about how to update or change the system images, see [Using System Utilities](#) on page 156.

Dual Image Configuration and Upgrade

Use the **Dual Image Configuration and Upgrade** feature to transfer a new firmware (code) image to the device, select which image to load during the next boot cycle, and add a description to each image on the device. The device uses the HTTP protocol to transfer the image, and the image is saved as the backup image.

To access this page, click **System > Firmware > Configuration and Upgrade** in the navigation menu.

Table 97: Dual Image Status Fields

Field	Description
Unit	Use this field to select the unit with the code image to activate, upgrade, delete, or describe.
Active	The active code file version. Use the icons to the right of the field to perform the file transfer. To transfer a new code image to the device, click the File Transfer icon. The Firmware Upgrade window opens. Click Choose File to browse to the file to transfer. After you select the appropriate file, click Begin Transfer to launch the HTTP transfer process. The active image is overwritten by the file that you transfer.

Table 97: Dual Image Status Fields (continued)

Field	Description
Backup	The backup code file version. Use the icons to the right of the field to perform the following tasks: <ul style="list-style-type: none"> To transfer a new code image to the device, click the File Transfer icon. The Firmware Upgrade window opens. Click Choose File to browse to the file to transfer. After you select the appropriate file, click Begin Transfer to launch the HTTP transfer process. If a backup image already exists on the device, it is overwritten by the file that you transfer. To delete the backup image from permanent storage, click the – (minus) icon. You must confirm the action before the image is deleted.
Next Active	Use this field to select the image version to load the next time this unit reboots.
Active Description	Use this field to specify a description to associate with the image that is currently the active code file.
Backup Description	Use this field to specify a description to associate with the image that is currently the backup code file.
Select File	Use this field to provide option to browse to the directory where the file is located and select the file to transfer to the device.
Digital Signature Verification	When this option is checked, the file download will be verified with the digital signature.
Status	Provides information about the status of the file transfer.

AutoInstall

The **AutoInstall** feature enables the configuration of a switch automatically when the device is turned on and, during the boot process, no configuration file is found in device storage. By communicating with a *DHCP* server, AutoInstall obtains an IP address for the switch and an IP address for a TFTP server. AutoInstall attempts to download a configuration file from the TFTP server and install in on the switch.

The DHCP server that the switch communicates with must provide the following information:

- The IP address and subnet mask (option 1) to be assigned to the switch.
- The IP address of a default gateway (option 3), if needed for IP communication.
- The identification of the TFTP server from which to obtain the boot file. This is given by any of the following fields, in the priority shown (highest to lowest):
 - The sname field of the DHCP reply.
 - The hostname of the TFTP server (option 66). Either the TFTP address or name—not both—is specified in most network configurations. If a TFTP hostname is given, a DNS server is required to translate the name to an IP address.
 - The IP address of the TFTP server (option 150).
 - The address of the TFTP server supplied in the siaddr field.
 - The name of the configuration file (boot file or option 67) to be downloaded from the TFTP server. **The boot file name must have a file type of *.cfg.**
- The IP addresses of DNS name servers (option 6). The IP addresses of DNS name servers should be returned from the DHCP server only if the DNS server is in the same LAN as the switch performing

AutoInstall. A DNS server is needed to resolve the IP address of the TFTP server if only the “sname” or option 66 values are returned to the switch.

After obtaining IP addresses for both the switch and the TFTP server, the AutoInstall feature attempts to download a host-specific configuration file using the boot file name specified by the DHCP server. If the switch fails to obtain the file, it will retry indefinitely.

To access this page, click **System > Firmware > AutoInstall**.

Table 98: AutoInstall Fields

Field	Definition
Admin Mode	The current administrative mode of the AutoInstall feature: <ul style="list-style-type: none"> • Start: AutoInstall is enabled, and the feature will attempt to automatically configure the device during the next boot cycle. • Stop: AutoInstall is disabled. The automatic process will begin only if no configuration file is located during the next boot cycle.
Persistent Mode	If this option is selected, the settings you configure on this page are automatically saved to persistent memory in the startup-config file when you apply the changes. If this option is not selected, the device treats these settings like any other applied changes: the changes are not retained across a reboot unless you save the configuration.
AutoSave Mode	If this option is selected, the downloaded configuration is automatically saved to persistent storage. If this option is not selected, you must explicitly save the downloaded configuration in non-volatile memory for the configuration to be available for the next reboot.
AutoReboot Mode	If this option is selected, the switch automatically reboots after a new image is successfully downloaded and makes the downloaded image the active image. If this option is not selected, the device continues to boot with the current image. The downloaded image will not become the active image until the device reboots.
Retry Count	When attempting to retrieve the DHCP-specified configuration file, this value represents the number of times the TFTP client on the device tries to use unicast requests before reverting to broadcast requests.
Status	The current status of the AutoInstall process.

Click **Refresh** to display the most recently configured AutoInstall state from the switch.

Managing Logs

The switch may generate messages in response to events, faults, or errors occurring on the platform as well as changes in configuration or other occurrences. These messages are stored both locally on the platform and forwarded to one or more centralized points of collection for monitoring purposes as well as long term archival storage. Local and remote configuration of the logging capability includes filtering of messages logged or forwarded based on severity and generating component.

The in-memory log stores messages in memory based upon the settings for message component and severity. On stackable systems, this log exists only on the management unit. Other platforms in the stack forward their messages to the management unit log. Access to in-memory logs on other than the management unit is not supported.

Log Configuration

The **Log Configuration** page allows administrators with the appropriate privilege level to configure the administrative mode and various settings for logging features on the switch.

To access this page, click **System > Logs > Configuration** in the navigation menu.

Table 99: Log Configuration Fields

Field	Description
Buffered Log Configuration	
Admin Mode	Enables or disables logging to the buffered (RAM) log file.
Behavior	Specifies what the device should do when the buffered log is full. It can either overwrite the oldest messages (Wrap) or stop writing new messages to the buffer (Stop on Full).
Command Logger Configuration	
Admin Mode	Enables or disables logging of the command-line interface (CLI) commands issued on the device.
Console Log Configuration	
Admin Mode	Enables or disables logging to any serial device attached to the host.
Severity Filter	Select the severity of the messages to be logged. All messages at and above the selected threshold are logged to the console. The severity can be one of the following: <ul style="list-style-type: none"> • emergency (0): The device is unusable. • alert (1): Action must be taken immediately. • critical (2): The device is experiencing primary system failures. • error (3): The device is experiencing non-urgent failures. • warning (4): The device is experiencing conditions that could lead to system errors if no action is taken. • notice (5): The device is experiencing normal but significant conditions. • info (6): The device is providing non-critical information. • debug (7): The device is providing debug-level information.
Persistent Log Configuration	
Admin Mode	Enable or disable logging to the persistent log. These messages are not deleted when the device reboots.
Severity Filter	Select the severity of the messages to be logged. All messages at and above the selected threshold are logged to the console. See the previous severity filter description for more information about each severity level.
Syslog Configuration	
Admin Mode	Enable or disable logging to configured syslog hosts. When the syslog admin mode is disabled the device does not relay logs to syslog hosts, and no messages will be sent to any collector/relay. When the syslog admin mode is enabled, messages will be sent to configured collectors/relays using the values configured for each collector/relay.
Protocol Version	The <i>RFC (Request for Comment)</i> version of the syslog protocol.
Local UDP Port	The <i>UDP (User Datagram Protocol)</i> port on the local host from which syslog messages are sent.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Buffered Log

The log messages the device generates in response to events, faults, errors, and configuration changes are stored locally on the device in the RAM (cache). This collection of log files is called the RAM log or buffered log. When the buffered log file reaches the configured maximum size, the oldest message is deleted from the RAM when a new message is added. If the system restarts, all messages are cleared.

To access the **Buffered Log** page, click **System > Logs > Buffered Log** in the navigation menu.

Table 100: Buffered Log Fields

Field	Description
Log Index	The position of the entry within the buffered log file. The most recent log message always has a Log Index value of 1.
Log Time	The time the entry was added to the log.
Severity	The severity level associated with the log entry. The severity can be one of the following: <ul style="list-style-type: none"> • emergency (0): The device is unusable. • alert (1): Action must be taken immediately. • critical (2): The device is experiencing primary system failures. • error (3): The device is experiencing non-urgent failures. • warning (4): The device is experiencing conditions that could lead to system errors if no action is taken. • notice (5): The device is experiencing normal but significant conditions. • info (6): The device is providing non-critical information. • debug (7): The device is providing debug-level information.
Component	The component that issued the log entry.
Description	The text description for the log entry.

Click **Refresh** to update the screen and associated messages.

Event Log

Use the **Event Log** page to display the event log, which is used to hold error messages for catastrophic events. After the event is logged and the updated log is saved in flash memory, the switch will reboot. The log can hold at least 2,000 entries (the actual number depends on the platform and OS), and is erased when an attempt is made to add an entry after it is full. The event log is preserved across system reboots.

To access this page, click **System > Logs > Event Log** in the navigation menu.

Table 101: Event Log Fields

Field	Description
Entry	The number of the entry within the event log. The most recent entry is first.
Type	The incident category that indicates the cause of the log entry: EVENT, ERROR, etc.
Filename	The 200 Series source code filename identifying the code that detected the event.
Line	The line number within the source file of the code that detected the event.
Task ID	The OS-assigned ID of the task reporting the event.
Code	The event code passed to the event log handler by the code reporting the event.
Time	The time the event occurred, measured from the previous reboot.

Click **Refresh** to update the screen and associated messages.

Hosts Log Configuration

Use the **Host Log Configuration** page to configure remote logging hosts where the switch can send logs.

To access this page, click **System > Logs > Hosts** in the navigation menu.

Table 102: Logging Hosts Fields

Field	Description
Host (IP Address/Host Name)	The IP address or DNS-resolvable host name of the remote host to receive log messages.
Status	Whether the host has been configured to be actively logging or not.
Port	The <u>UDP</u> port on the logging host to which syslog messages are sent.
Severity Filter	Severity level threshold for log messages. All log messages with a severity level at and above the configured level are forwarded to the logging host.
Transport Mode	Transport mode used while sending messages to syslog servers. Supported modes are UDP and <u>TLS</u> . If TLS is not configured, default transport mode is UDP.
Authentication Mode	Using TLS security user can configure anonymous authentication mode, in which no client authentication is done by the syslog server. For x509/name authentication mode, two-way authentication is done both by syslog client and client authentication by syslog server side.
Certificate Index	The index used for identifying corresponding certificate files.

Use the buttons to perform the following tasks:

- To add a logging host, click **Add** and configure the desired settings.
- To change information for an existing logging host, select the checkbox for the entry and click **Edit**. You cannot edit the host name or address of a host that has been added.
- To delete a configured logging host from the list, select the checkbox for each entry to delete and click **Remove**.

After you add a logging host, the screen displays additional fields.

Table 103: Host Log Configuration Fields

Field	Description
IP Address/Host Name	The IP address or DNS-resolvable host name of the remote host to receive log messages.
Transport Mode	Transport mode used while sending messages to syslog servers. Supported modes are UDP and TLS. If TLS is not configured then default transport mode is UDP.
Authentication Mode	Using TLS security user can configure anonymous authentication mode, in which no client authentication is done by the syslog server. For x509/name authentication mode, two way authentication is done both by syslog client and client authentication by syslog server side.
Certificate Index	The index used for identifying corresponding certificate files.
Port	The UDP port on the logging host to which syslog messages are sent.
Severity Filter	Severity level threshold for log messages. All log messages with a severity level at and above the configured level are forwarded to the logging host.

Adding a Remote Logging Host

Use the following procedures to add, configure, or delete a remote logging host.

- 1 From the **Host** field, select **Add** to add a new host, or select the IP address of an existing host to configure the host.
If you are adding a new host, enter the IP address of the host in the **IP Address** field and click **Submit**. The screen refreshes, and additional fields appear.
- 2 In the **Port** field, type the port number on the remote host to which logs should be sent.
- 3 Select the severity level of the logs to send to the remote host.
- 4 Click **Submit** to apply the changes to the system.

Deleting a Remote Logging Host

To delete a remote logging host from the configured list, select the IP address of the host from the Host field, and then click **Delete**.

Syslog Source Interface Configuration

Use the **Syslog Source Interface Configuration** page to specify the physical or logical interface to use as the logging (*syslog*) client source interface. When an IP address is configured on the source interface, this address is used for all Syslog communications between the local logging client and the remote Syslog server. The IP address of the designated source interface is used in the IP header of Syslog management protocol packets. This allows security devices, such as firewalls, to identify all source packets coming from a specific device.

To access this page, click **System > Logs > Source Interface Configuration** in the navigation menu.

Table 104: Syslog Source Interface Configuration Fields

Field	Description
Type	The type of interface to use as the source interface: <ul style="list-style-type: none"> • None: The primary IP address of the originating (outbound) interface is used as the source address. • Interface: The primary IP address of a physical port is used as the source address. • Loopback: The primary IP address of the loopback interface is used as the source address. A loopback is always reachable, as long as any routing interface is up. • VLAN: The primary IP address of a VLAN routing interface is used as the source address.
Interface	When the selected Type is Interface, select the physical port to use as the source interface.
Loopback Interface	When the selected Type is Loopback, select the loopback interface to use as the source interface.
VLAN ID	When the selected Type is VLAN, select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Persistent Log

Use the **Persistent Log** page to view the persistent log messages.

To access this page, click **System > Log > Persistent Log** in the navigation menu.

Table 105: Persistent Log Fields

Field	Description
Log Index	The position of the entry within the buffered log file. The most recent log message always has a Log Index value of 1.
Log Time	The time the entry was added to the log.
Severity	The severity level associated with the log entry. The severity can be one of the following: <ul style="list-style-type: none"> • emergency (0): The device is unusable. • alert (1): Action must be taken immediately. • critical (2): The device is experiencing primary system failures. • error (3): The device is experiencing non-urgent failures. • warning (4): The device is experiencing conditions that could lead to system errors if no action is taken. • notice (5): The device is experiencing normal but significant conditions. • info (6): The device is providing non-critical information. • debug (7): The device is providing debug-level information.

Table 105: Persistent Log Fields (continued)

Field	Description
Component	The component that has issued the log entry.
Description	The text description for the log entry.

Configuring and Searching the Forwarding Database

The *FDB* maintains a list of MAC addresses after having received a packet from this MAC address. The transparent bridging function uses the forwarding database entries to determine how to forward a received frame.

Switch Configuration

Use the **Switch Configuration** page to set the amount of time to keep a learned MAC address entry in the forwarding database. The forwarding database contains static entries, which are never aged out, and dynamically learned entries, which are removed if they are not updated within a given time.

To access this page, click **System > Basic Configuration > Switch** in the navigation menu.

Table 106: Switch Configuration Fields

Field	Description
802.3x Flow Control Mode	Enable or disable 802.3x flow control on the switch. IEEE 802.3x flow control works by pausing a port when the port becomes oversubscribed. It also allows a port to drop all traffic for small bursts of time during the congestion condition. This can lead to high-priority and/or network control traffic loss. When enabled, flow control allows lower speed switches to communicate with higher-speed switches by requesting that the higher-speed switch refrain from sending packets. Transmissions are temporarily halted to prevent buffer overflows.
MAC Address Aging Interval	The MAC address table (forwarding database) contains static entries, which never age out, and dynamically-learned entries, which are removed if they are not updated within a given time. Specify the number of seconds a dynamic address should remain in the MAC address table after it has been learned.



Note

IEEE 802.1D recommends a default of 300 seconds, which is the factory default.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Configuring Power Over Ethernet (PoE) and PoE Statistics

Use these pages to view *PoE (Power over Ethernet)* status information, configure global PoE settings, configure PoE settings on interfaces and view PoE interface statistical information.

PoE Configuration

Use the **PoE Configuration** page to view *PoE* status information and configure global PoE settings.

To access this page, click **System > PoE > Configuration** in the navigation menu.

Table 107: PoE Configuration Fields

Field	Description
Firmware Version	The firmware version of the PoE software component.
Operational Status	The current status of the switch PoE functionality, which can be one of the following: <ul style="list-style-type: none"> • On: At least one port on the switch is delivering power to a connected device. • Off: The PoE functionality is operational but no ports are delivering power. • Faulty: The PoE functionality is not operational.
Total Power Available	The total power in mWatts that can be provided by the switch.
Threshold Power	When the PoE power being used exceeds this threshold, a trap is generated to the system log to alert the system administrator of high power usage. This value is determined by the configurable System Usage Threshold percentage.
Consumed Power	The amount of power in mWatts currently being consumed by connected PoE devices.
System Usage Threshold	A percentage of the total power available. This percentage determines the threshold power.
Power Management Mode	The method by which the PoE controller determines supplied power, which can be one of the following: <ul style="list-style-type: none"> • Static: The power allocated to each port is reserved and is not available to any other port, even when less than the maximum allocation is being used. • Dynamic: The power allocated to each port is not reserved. Unused power may be allocated from one port to another as needed, up to the power limit defined for each port.
Port Auto Reset Mode	When enabled, the switch automatically resets a PoE port if an error condition occurs. When disabled, the administrator must reset the port manually.
Traps	When enabled, <i>SNMP</i> traps will be generated when certain events occur. Trap events include a change in whether power is being delivered on a port and when the power usage threshold is exceeded.

- If you make any changes to the page, click **Submit** to apply the changes to the system.
- Click **Refresh** to redisplay the page with the current data from the switch.

PoE Port Configuration

Use the **PoE Port Configuration** page to configure *PoE* settings on interfaces.

To access this page, click **System > PoE > Port Configuration** in the navigation menu.

Table 108: PoE Port Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When configuring PoE settings, this field identifies the interface(s) being configured.
Admin Mode	Whether PoE is administratively enabled or disabled on the interface.
Priority	The priority of the port when allocating available power. Power is delivered to the higher priority ports when needed before providing it to the lower priority ports. Possible values are Critical, High, Medium, and Low.
High Power Mode	When enabled, the port supports the PoE+ power standard, which allows for providing up to 30W of power. When disabled, the port supports the original PoE standard only, which allows for providing up to 15.4W of power.
Power Limit Type	The type of power limiting used for the port, which can be one of the following: <ul style="list-style-type: none"> • Class: The device class determines the power limit. The switch learns the class of the device through the receipt of <i>LLDP (Link Layer Discovery Protocol)</i> messages. • User: The power limit is user defined, overriding the LLDP information. When set, the Power Limit field is enabled.
Power Limit	The power limit for the port, which can be specified. This field displays only when Power Limit Type is set to User.
Detection Type	The protocol(s) that can be used to detect the presence of a PD when connected to a PoE port. The IEEE specification 802.3af (Dot3af) specifies various detection algorithms. Some PDs use legacy detection algorithms that were in place prior to the 802.3af standard, which can be one of the following: <ul style="list-style-type: none"> • Legacy: The switch uses a legacy detection scheme not defined in 802.3af. • 4Pt-Dot3af: The switch uses the 802.3af 4-point detection scheme only. • 4Pt-Dot3af + Legacy: The switch uses the 802.3af 4-point detection scheme, followed by the legacy detection scheme. • 2Pt-Dot3af : The switch uses the 802.3af 2-point detection scheme. • 2Pt-Dot3af + Legacy: The switch uses the 802.3af 2-point detection scheme, followed by the legacy detection scheme. • None: No detection is performed.
Timer Schedule	The time range from the list of time ranges configured on the system.

Table 108: PoE Port Configuration Fields (continued)

Field	Description
Status	The status of the port as a provider of PoE. Such devices are referred to as PSE. The status can be one of the following: <ul style="list-style-type: none"> • Disabled: The PSE is disabled. • Delivering Power: The PSE is delivering power. • Fault: The PSE has experienced a fault condition. • Test: The PSE is in test mode. • Other Fault: The PSE has experienced a variable error condition. • Searching: The PSE is transitioning between states. • Requesting Power: The PSE is currently not able to deliver power because power is unavailable to the port.
Fault Status	The error when PSE port is in fault status, which can be one of the following: <ul style="list-style-type: none"> • None: PSE port is not in any error state. • MPS Absent: PSE port has detected absence of main power supply. • Short: PSE port has detected a short circuit condition. • Overload: PD connected to PSE port tried to draw more power than permissible by the hardware. • Power Denied: PSE port has been denied power due to administrative action or shortage of power.

To display additional PoE interface information, select an entry and click **Details**. The following information describes the fields in the **Details** window.

Table 109: PoE Port Entry Details Fields

Field	Description
High Power	Whether high power mode is enabled or disabled.
Max Power	If Power Limit Type for the port is set to User (user defined), this field displays the configured power limit. If Power Limit Type is set to Class, this field is blank.
Class	If Power Limit Type is set to Class, this field displays the class of the connected device, as learned in LLDP messages. Possible values are Unknown and Class 0 through Class 4. A higher class value indicates that the device requires higher power.
Output Voltage	The voltage being applied to the connected device.
Output Current	The current in milliamps being drawn by the powered device.
Output Power	The power in mWatts being drawn by the connected device.
Temperature	The temperature measured at the PoE port.

Use the buttons to perform the following tasks:

- To configure the settings for one or more interfaces, select each entry to modify and click **Edit**.
- To apply the same settings to all interfaces, click **Edit All**.

PoE Port Statistics

Use the **PoE Port Statistics** page to view [PoE](#) interface statistical information.

To access this page, click **System > PoE > Statistics** in the navigation menu.

Table 110: PoE Port Statistics Fields

Field	Description
Interface	The interface associated with the rest of the data on the page.
Overload Counter	Number of times there has been a power overload. Power overload occurs when a powered device connected to a port tries to draw more power than permissible by the hardware.
Short Counter	Number of times there has been a short circuit condition.
Power Denied Counter	Number of times the powered device has been denied power. Power is denied due to administrative action or shortage of power.
MPS Absent Counter	Number of times power has stopped because the powered device was not detected.
Invalid Signature Counter	Number of times an invalid signature was received. Signature detection is a stage in detecting the presence of a powered device, where a resistance value on the powered device is expected to be found within a particular range.

Click **Refresh** to redisplay the page with the current data from the switch.

Viewing Device Port Information

The pages in the Port folder allow you to view and monitor the physical port information for the ports available on the switch.

Port Summary

Use the **Port Summary** page to view the settings for all physical ports on the platform.

To access this page, click **System > Port > Summary** in the navigation menu.

Table 111: Port Summary Fields

Field	Description
Interface	Identifies the port that the information in the rest of the row is associated with.
Interface Index	The interface index object value assigned by the IF-MIB. This value is used to identify the interface when managing the device by using SNMP .

Table 111: Port Summary Fields (continued)

Field	Description
Type	<p>For most ports this field is blank. Otherwise, the possible values are:</p> <ul style="list-style-type: none"> • Normal: The port is a normal port, which means it is not a LAG member or configured for port mirroring. • Trunk Member: The port is a member of a LAG. • Mirrored: Indicates that the port has been configured as a monitoring port and is the source port in a port mirroring session. For more information about port monitoring and probe ports, see Mirroring on page 129. • Probe: Indicates that the port has been configured as a monitoring port and is the destination port in a port mirroring session. For more information about port monitoring and probe ports, see Mirroring on page 129.
Admin Mode	<p>Shows the port control administration state, which can be one of the following:</p> <ul style="list-style-type: none"> • Enabled: The port can participate in the network (default). • Disabled: The port is administratively down and does not participate in the network.
Physical Mode	<p>Shows the speed and duplex mode at which the port is configured:</p> <ul style="list-style-type: none"> • Auto: The duplex mode and speed will be set by the auto-negotiation process. The port's maximum capability (full duplex and 100 Mbps) will be advertised. • <Speed> Half Duplex: The port speeds available from the menu depend on the platform on which the 200 Series software is running and which port you select. In half-duplex mode, the transmissions are one-way. In other words, the port does not send and receive traffic at the same time. • <Speed> Full Duplex: The port speeds available from the menu depend on the platform on which the 200 Series software is running and which port you select. In half-duplex mode, the transmissions are two-way. In other words, the port can send and receive traffic at the same time.
Physical Status	Indicates the port speed and duplex mode at which the port is operating.
STP Mode	<p>The STP Administrative Mode associated with the port or LAG. STP is a Layer 2 protocol that provides a tree topology for switches on a bridged LAN. STP allows a network to have redundant paths without the risk of network loops, by providing a single path between end stations on a network. The possible values for STP mode are:</p> <ul style="list-style-type: none"> • Enable: Spanning tree is enabled for this port. • Disable: Spanning tree is disabled for this port.
LACP Mode	<p>The LACP (Link Aggregation Control Protocol) administration state. The mode must be enabled in order for the port to participate in Link Aggregation. This field can have the following values:</p> <ul style="list-style-type: none"> • Enable: Specifies that the port is allowed to participate in a port channel (LAG), which is the default mode. • Disable: Specifies that the port cannot participate in a port channel (LAG). • N/A: For LAG ports.
Link Status	Whether the Link is up or down.
The following fields can be accessed by selecting a port and clicking Edit :	

Table 111: Port Summary Fields (continued)

Field	Description
Link Trap	<p>This object determines whether or not to send a trap when link status changes. The factory default is enabled.</p> <ul style="list-style-type: none"> • Enable: Specifies that the system sends a trap when the link status changes. • Disable: Specifies that the system does not send a trap when the link status changes.
Maximum Frame Size	The maximum Ethernet frame size the interface supports or is configured to support. The maximum frame size includes the Ethernet header, CRC, and payload.
Broadcast Storm Recovery Mode	<p>Specifies the broadcast storm control threshold for the port. Broadcast storm control limits the amount of broadcast frames accepted and forwarded by the port. If the broadcast traffic on the Ethernet port exceeds the configured threshold, the system blocks (discards) the broadcast traffic.</p> <p>Specifies the broadcast storm recovery action to either Shutdown or Trap for specific interface. If configured to Shutdown, the interface which receives broadcast packets at a rate which is above threshold is diagnostically disabled. The Trap option sends trap messages at approximately every 30 seconds until broadcast storm control recovers.</p>
Multicast Storm Recovery Level	<p>Specifies the multicast storm control threshold for the port. Multicast storm control limits the amount of multicast frames accepted and forwarded by the port. If the multicast traffic on the Ethernet port exceeds the configured threshold, the system blocks (discards) the multicast traffic.</p> <p>Specifies the multicast storm recovery action to either Shutdown or Trap for specific interface. If configured to Shutdown, the interface which receives multicast packets at a rate which is above threshold is diagnostically disabled. The Trap option sends trap messages at approximately every 30 seconds until multicast storm control recovers.</p>
Unicast Storm Recovery Level	<p>Specifies the unicast storm control threshold for the port. Unicast storm control limits the amount of unicast frames accepted and forwarded by the switch. If the unicast traffic on the Ethernet port exceeds the configured threshold, the system blocks (discards) the unicast traffic.</p> <p>Specifies the unicast storm recovery action to either Shutdown or Trap for specific interface. If configured to Shutdown, the interface which receives unicast packets at a rate which is above threshold is diagnostically disabled. The Trap option sends trap messages at approximately every 30 seconds until unicast storm control recovers.</p>

Click **Refresh** to redisplay the most current information from the router.

Port Description

Use the **Port Description** page to configure a human-readable description of the port.

To access this page, click **System > Port > Description** in the navigation menu.

Table 112: Port Description Fields

Field	Description
Interface	Select the interface for which data is to be displayed or configured.
Port Description (Input field)	A user-configurable description to help identify the port. To add a description to a port, select the port or <i>LAG</i> from the Interface drop-down menu, type a description in the Port Description field, and click Submit .
Physical Address	Displays the physical address of the specified interface.
PortList Bit Offset	Displays the bit offset value which corresponds to the port when the MIB object type PortList is used to manage in <i>SNMP</i> .
Interface Index	The interface index object value assigned by the IF-MIB. This value is used to identify the interface when managing the device by using SNMP.
Port Description	Shows the configured port description. By default, the port does not have an associated description.

- If you change a port description, click **Submit** to apply the change to the system.
- Click **Refresh** to redisplay the page with the latest information from the router.

Cable Test

The **Port Cable Test** feature enables you to determine the cable connection status on a selected port. You can also obtain an estimate of the length of the cable connected to the port, if the *PHY (Physical Interface Transceiver)* on the ports supports this functionality.



Note

The cable test feature is supported only for copper cable. It is not supported for optical fiber cable.

To access this page, click **System > Port > Cable Test**.

The page displays with additional fields when you click **Test Cable**. The fields that display depend on the cable test results.

Table 113: Cable Test Fields

Field	Description
Interface	If the test has not been performed, this is the only field that displays. Select the interface to test. After the test has been performed, this field shows the interface that was tested.
Cable Status	This displays the cable status as Normal, Open, or Short. <ul style="list-style-type: none"> • Normal: The cable is working correctly. • Open: The cable is disconnected or there is a faulty connector. • Open and Short: There is an electrical short in the cable. • Cable Status Test Failed: The cable status could not be determined. The cable may in fact be working. This field is displayed after you click Test Cable and results are available.

Table 113: Cable Test Fields (continued)

Field	Description
Cable Length	<p>The estimated length of the cable. If the cable length cannot be determined, Unknown is displayed. This field shows the range between the shortest estimated length and the longest estimated length.</p> <p>Note: This field displays a value only when the Cable Status is Normal; otherwise, this field is blank.</p>
Failure Location	<p>The estimated distance from the end of the cable to the failure location.</p> <p>Note: This field displays a value only when the Cable Status is Open or Short; otherwise, this field is blank.</p>

Select a port and click **Test Cable** to display its status.

If the port has an active link while the cable test is run, the link can go down for the duration of the test. The test may take several seconds to run.

The command returns a cable length estimate if this feature is supported by the PHY for the current link speed. Note that if the link is down and a cable is attached to a 10/100 Ethernet adapter, then the cable status may display as Open or Short because some Ethernet adapters leave unused wire pairs unterminated or grounded.

Mirroring

Port mirroring selects the network traffic for analysis by a network analyzer. This is done for specific ports of the switch. As such, many switch ports are configured as source ports and one switch port is configured as a destination port. You can configure how traffic is mirrored on a source port. Packets that are received on the source port, that are transmitted on a port, or are both received and transmitted, can be mirrored to the destination port.

The packet that is copied to the destination port is in the same format as the original packet on the wire. This means that if the mirror is copying a received packet, the copied packet is VLAN tagged or untagged as it was received on the source port. If the mirror is copying a transmitted packet, the copied packet is VLAN tagged or untagged as it is being transmitted on the source port.

Use the **Multiple Port Mirroring** page to define port mirroring sessions.

To access this page, click **System > Port > Mirroring** in the navigation menu.

Use the buttons to perform the following tasks:

- To configure the administrative mode for a port mirroring session, click **Configure Session** and configure the desired settings.
- To configure destination as Remote VLAN or probe port, click **Edit** and configure the desired settings.

- To configure source as Remote VLAN or VLAN or one or more source ports for the mirroring session and to determine which traffic is mirrored (Tx, Rx, or both), click **Configure Source** and configure the desired settings.
- To remove one or more source ports from the port mirroring session, select the checkbox for each source port to remove and click **Remove Source**.

Table 114: Multiple Port Mirroring Fields

Field	Description
Session ID	The port mirroring session ID. The number of sessions allowed is platform specific.
Mode	The administrative mode for the selected port mirroring session. If the mode is disabled, the configured source is not mirroring traffic to the destination.
Destination	<ul style="list-style-type: none"> • Remote VLAN: At the source switch the destination can be configured as RSPAN VLAN with reflector-port. The reflector-port forwards the mirrored traffic towards the destination switch. This port has to be configured with RSPAN VLAN membership. • Interface: If port configured as a interface or probe port. This port receives traffic from all configured source ports. • None: The destination is not configured.
IP ACL	The IP access-list ID or name attached to the port mirroring session.
MAC ACL	The MAC access-list name attached to the port mirroring session.
Source	The ports or VLAN configured to mirror traffic to the destination. You can configure multiple source ports or one source VLAN per session. The source VLAN can also be a remote VLAN.
Direction	<p>The direction of traffic on the source port(s) that is sent to the probe port. Possible values are:</p> <ul style="list-style-type: none"> • Tx and Rx: Both ingress and egress traffic. • Rx: Ingress traffic only. • Tx: Egress traffic only.

Configuring a Port Mirroring Session

Note



If an interface participates in some VLAN and is a LAG member, this VLAN cannot be assigned as a source VLAN for a Monitor session. At the same time, if an interface participates in some VLAN and this VLAN is assigned as a source VLAN for a Monitor session, the interface can be assigned as a LAG member.

- 1 From the **Multiple Port Mirroring** page, click **Configure Session** to display the **Session Configuration** page.

- 2 Configure the following fields:

Table 115: Multiple Port Mirroring—Destination Configuration

Field	Description
Session ID	The port mirroring session ID. The number of sessions allowed is platform specific.
Mode	The administrative mode for the selected port mirroring session. If the mode is disabled, the configured source is not mirroring traffic to the destination.
IP ACL	The ID of the IP ACL to apply to traffic from the source. Only traffic that matches the rules in the ACL is mirrored to the destination.
MAC ACL	The ID of the MAC ACL to apply to traffic from the source. Only traffic that matches the rules in the ACL is mirrored to the destination.

- 3 Click **Submit** to apply the changes to the system.

Configuring a Port Mirroring Source

Note



If an interface participates in some VLAN and is a [LAG](#) member, this VLAN cannot be assigned as a source VLAN for a Monitor session. At the same time, if an interface participates in some VLAN and this VLAN is assigned as a source VLAN for a Monitor session, the interface can be assigned as a LAG member.

- 1 From the **Multiple Port Mirroring** page, click **Configure Source** to display the **Source Configuration** page.
- 2 Configure the following fields:

Table 116: Multiple Port Mirroring—Source Configuration

Field	Description
Session ID	The port mirroring session ID. The number of sessions allowed is platform specific.
Type	The type of interface to use as the source: <ul style="list-style-type: none"> • None: The source is not configured. • Remote VLAN: The VLAN configured as the RSPAN VLAN is the source. In an RSPAN configuration, the remote VLAN is the source on the destination device that has a physical port connected to the network traffic analyzer. • VLAN: Traffic to and from a configured VLAN is mirrored. In other words, all the packets sent and received on all the physical ports that are members of the VLAN are mirrored. • Interface: Traffic is mirrored from one or more physical ports on the device.
Remote VLAN	The VLAN that is configured as the RSPAN VLAN.
VLAN ID	The VLAN to use as the source. Traffic from all physical ports that are members of this VLAN is mirrored. This field is available only when the selected Type is VLAN.

Table 116: Multiple Port Mirroring—Source Configuration (continued)

Field	Description
Available Source port(s)	The physical port or ports to use as the source. To select multiple ports, hold down [Ctrl] and click each port. This field is available only when the selected Type is Interface.
Direction	The direction of traffic on the source port (or source ports) or VLAN that is sent to the specified destination. A source VLAN mirrors all received and transmitted packets to the destination. Possible values for source ports are: <ul style="list-style-type: none"> • Tx and Rx: Both ingress and egress traffic • Rx: Ingress traffic only • Tx: Egress traffic only

- 3 Click **Submit** to apply the changes to the system.

Adding a Port Mirroring Session



Note

A port will be removed from a VLAN or *LAG* when it becomes a destination mirror.

- 1 From the **Port Mirroring** page, click **Add** to display the **Add Source Ports** page.
- 2 Configure the following fields:

Table 117: Multiple Port Mirroring—Add Source Ports Fields

Field	Description
Session	Specifies the monitoring session.
Source Port	Select the unit and port from which traffic is mirrored. Up to eight source ports can be mirrored to a destination port.
Direction	Select the type traffic monitored on the source port, which can be one of the following: <ul style="list-style-type: none"> • Tx and Rx: Monitors transmitted and received packets • Rx: Monitors received packets only • Tx: Monitors transmitted packets only

- 3 Click **Add** to apply the changes to the system.

The new port mirroring session is enabled for the unit and port, and the device is updated. The source port appears in the **Source Port** list on the **Multiple Port Mirroring** page.

Removing or Modifying a Port Mirroring Session

The source ports are removed from the port mirroring session, and the device is updated.

- 1 From the **Port Mirroring** page, click **Remove Source Port**.
- 2 Select one or more source ports to remove from the session.
Use the **[Ctrl]** key to select multiple ports to remove.
- 3 Click **Remove**.

Mirroring Summary

Use the **Multiple Port Mirroring Summary** page to view the port mirroring summary.

To access this page, click **System > Port > Mirroring Summary** in the navigation menu.

Table 118: Multiple Port Mirroring Summary Fields

Field	Description
Session ID	The port mirroring session ID. The number of sessions allowed is platform specific.
Admin Mode	The administrative mode for the selected port mirroring session. If the mode is disabled, the configured source is not mirroring traffic to the destination.
Probe Port	The interface that receives traffic from all configured source ports.
Src VLAN	The VLAN configured to mirror traffic to the destination. You can configure one source VLAN per session. The source VLAN can also be a remote VLAN.
Mirrored Port	The ports configured to mirror traffic to the destination. You can configure multiple source ports per session.
Reflector Port	This port carries all the mirrored traffic at source switch.
Src RVLAN	The VLAN configured as the RSPAN VLAN is the source. In an RSPAN configuration, the remote VLAN is the source on the destination device that has a physical port connected to the network traffic analyzer.
Dst RVLAN	Traffic is mirrored to the VLAN on the system that is configured as the RSPAN VLAN. In an RSPAN configuration, the destination should be the Remote VLAN on any device that does not have a port connected to the network traffic analyzer.
Type	The type of traffic on the source port (or source ports) or VLAN that is sent to the specified destination. A source VLAN mirrors all received and transmitted packets to the destination. Possible values for source ports are: <ul style="list-style-type: none"> • Tx and Rx: Both ingress and egress traffic • Rx: Ingress traffic only • Tx: Egress traffic only
IP ACL	The ID of the IP ACL to apply to traffic from the source. Only traffic that matches the rules in the ACL is mirrored to the destination.
MAC ACL	The ID of the MAC ACL to apply to traffic from the source. Only traffic that matches the rules in the ACL is mirrored to the destination.

Click **Refresh** to redisplay the page with the latest information from the router.

Green Ethernet Status

For platforms that include Green Energy features, the **Green Ethernet Status** page shows status information about the Green Ethernet feature on the device. The Green Ethernet feature is designed to reduce per-port power usage.

To access this page, click **System > Advanced Configuration > Green Ethernet > Status** in the navigation menu.

Table 119: Green Ethernet Status Fields

Field	Description
Estimated Energy Savings	The estimated cumulative energy saved on the device (Watts times hours) because of the Green Ethernet feature.
Estimated Power Savings	The estimated percentage of power saved on all ports. For example, 10% means that the device required 10% less power than it would have required if the Green Ethernet feature had not been present.
Current Power Consumption	The estimated power consumption by all ports.
Unit	The device's Unit ID.
Energy-Detect	Whether Energy Detect mode is present on the device. When the Energy Detect mode is enabled and a port link is down, the <i>PHY</i> automatically goes down for a short period of time and then wakes up to check link pulses. This mode reduces power consumption on the port when no link partner is present.
Short-Reach	Whether the Short-Reach cable mode is present on the device. When present and enabled, short-reach cable mode performs a cable test when the port link is up. If the cable that connects the port to its link partner has a length less than 10m, PHYs are placed in low-power mode (nominal power).
EEE	Whether EEE (Energy Efficient Ethernet) is present on the device. EEE enables ports to enter a low-power mode to reduce power consumption during periods of low link utilization. EEE is defined by IEEE 802.3az. EEE enables both the send and receive sides of the link to disable some functionality for power savings when the link is lightly loaded.
LPI-History	Whether the device is able to provide historical data about the amount of time it has spent in LPI (low-power idle) mode.
LLDP-Cap-Exchg	Whether the device is able to exchange information about its power capabilities with link partners by transmitting information in <i>LLDP</i> data units.
Pwr-Usg-Est	Whether the device is able to provide estimates of its power consumption.

Green Ethernet Configuration

For platforms that include Green Energy features, the **Green Ethernet Configuration** page is used to configure the global settings for the EEE (Energy Efficient Ethernet) settings on the device.

To access this page, click **System > Advanced Configuration > Green Ethernet > Configuration** in the navigation menu.

Table 120: Green Ethernet Configuration Fields

Field	Description
EEE LPI History Sampling Interval	The amount of time to wait between collecting LPI (Low-Power Idle) samples on the device.
EEE LPI History Maximum	The number of LPI samples to store in the buffer.
EEE Low Power Idle	The administrative mode of EEE LPI on the device. When enabled, the ports can enter a low-power mode to reduce power consumption during periods of low link utilization.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Green Ethernet Interface Configuration

For platforms that include Green Energy features, the **Green Ethernet Interface Configuration** page is used to configure per-port Green Ethernet settings. Only interfaces that are capable of supporting Green Ethernet modes appear in the table. To configure the settings for one or more interfaces, select each interface to configure and click Edit. The same settings are applied to all selected interfaces.

To access this page, click **System > Advanced Configuration > Green Ethernet > Interface** in the navigation menu.

Table 121: Green Ethernet Interface Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When configuring the settings for one or more interfaces, this field identifies the interfaces being configured.
Energy Detection	The administrative status of Energy Detect mode on the interface. When the Energy Detect mode is enabled and a port link is down, the <i>PHY</i> automatically goes down for short period of time and then wakes up to check link pulses. This mode reduces power consumption on the port when no link partner is present.
Energy Detection Status	The current operational state of Energy Detect mode – either Active or Inactive.
EEE Low Power Idle	The administrative mode of LPI (Low-Power Idle) on the interface. LPI can reduce power consumption on the interface during periods where no traffic is present on the interface. Enabling this mode does not affect link status and should not cause traffic loss. Note that LPI mode is available only if the interface Physical Mode is Auto Negotiate.
EEE Idle Time	The amount of time allowed for the interface to move to an LPI state.
EEE Wake Time	The system wake time that the interface transmits when it is enabled for EEE (Energy Efficient Ethernet). The wake time is the amount of time allowed to wake up from the low-power state that occurs when no data is transmitted.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Green Ethernet Local Interface Status

For platforms that include Green Energy features, the **Green Ethernet Local Interface Status** page shows information that each EEE (Energy Efficient Ethernet)-enabled interface transmits in the *LLDP* TLVs (type length values) to its link partner (the remote system). The TLVs are defined in the IEEE 802.1AB standard and provide information about the capabilities of the local device.

To access this page, click **System > Advanced Configuration > Green Ethernet > Local** in the navigation menu.

Table 122: Green Ethernet Local Interface Status Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. The table displays all interfaces that are enabled for EEE.
Tw_sys_tx	The system wake time (Tw_sys) that the interface transmits. The wake time is the amount of time allowed to wake up from the low-power state that occurs when no data is transmitted.
Tw_sys_tx Echo	The system wake time the interface sends to the link partner when it receives a Tw_sys_tx request from the link partner.
Tw_sys_rx	The system wake time that the local interface requests from the remote link partner.
Tw_sys_rx Echo	The remote system's receive Tw_sys that was used by the local system to compute the Tw_sys that it can support.
Fallback Tw_sys	The value of fallback Tw_sys that the local system requests from the remote system. The fallback is the second preference of the receiving system when requesting the Tw_sys from its transmitting partner.
Tx DLL Enabled	The initialization status of the EEE transmit DLL (Data Link Layer) management function on the local system.
Tx DLL Ready	The DLL ready transmission status of the interface. This field indicates whether the transmission system initialization is complete and is ready to update/transmit LLDP Data Units (LLDPDUs) containing the EEE TLVs.
Rx DLL Enabled	The status of the EEE capability negotiation on the local interface.
Rx DLL Ready	The DLL ready receive status of the interface. This field indicates whether the local interface initialization is complete and is ready to update/receive LLDPDUs containing EEE TLVs.

Green Ethernet Remote Device Status

For platforms that include Green Energy features, the **Green Ethernet Remote Device Status** page shows the information that each EEE (Energy Efficient Ethernet)-enabled interface transmits in the *LLDP* TLVs (type length values) to its link partner (the remote system). The TLVs are defined in the IEEE 802.1AB standard and provide information about the capabilities of the local device.

To access this page, click **System > Advanced Configuration > Green Ethernet > Remote Devices** in the navigation menu.

Table 123: Green Ethernet Remote Device Status Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. The table displays all interfaces that are enabled for EEE and have received EEE TLVs from a link partner.
Tw_sys_tx	The system wake time (Tw_sys) that the interface received from its link partner.
Tw_sys_tx Echo	The system wake time the remote system sends to the local interface when it receives a Tw_sys_tx request from the local interface.

Table 123: Green Ethernet Remote Device Status Fields (continued)

Field	Description
Tw_sys_rx	The system wake time that the remote link partner requests from the local interface.
Tw_sys_rx Echo	The local system's receive Tw_sys used by the remote system to compute the Tw_sys that it can support.
Fallback Tw_sys	The value of fallback Tw_sys that the remote system requests from the local system. The fallback is the second preference of the receiving system when requesting the Tw_Sys from its transmitting partner.

Green Ethernet Statistics

For platforms that include Green Energy features, the **Green Ethernet Statistics** page shows information about the amount of energy saved for each port. This page also allows you to enable or disable the green mode features that the switch supports. The green mode features can be controlled on a port-by-port basis.

To access this page, click **System > Advanced Configuration > Green Ethernet > Statistics** in the navigation menu.

Use the buttons to perform the following tasks:

- **Clear** resets all Green Ethernet statistics counters on this page to 0.
- **Refresh** refreshes the data on the screen with the present state of the data in the switch.

Table 124: Green Ethernet Statistics Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. The table includes all interfaces that are enabled for EEE.
Rx Low Power Idle Event Count	The number of times the local interface has entered a low-power idle state.
Rx Low Power Idle Duration	The amount of time (in 10 microsecond increments) the local interface has spent in a low-power idle state.
Tx Low Power Idle Event Count	The number of times the link partner has entered a low-power idle state.
Tx Low Power Idle Duration	The amount of time (in 10 microsecond increments) the link partner has spent in a low-power idle state.
Time Since Counters Last Cleared	The amount of time since the statistics on this page were reset to zero.

Green Ethernet EEE Interface History

Use the **Green Ethernet EEE Interface History** page to set the sampling interval for EEE LPI data and to specify the number of samples to keep. From this page, you can also view per-port EEE LPI data.

To access this page, click **System > Advanced Configuration > Green Ethernet > EEE History** in the navigation menu.

Table 125: Green Ethernet EEE Interface History Fields

Field	Description
Interface	Select the interface with the green mode information to view or configure.
Sample No.	A unique number that identifies the sample.
Age	The amount of time that has passed since the sample was recorded.
% Time in LPI since last sample	The percentage of time the interface has spent in LPI mode since the last sample was taken.
% Time in LPI since last reset	The percentage of time the interface has spent in LPI mode since the last time the EEE statistics were cleared.

Configuring and Viewing Device Slot Information

The pages in the Slot folder provide information about the cards installed in the slots on the switch. The physical location of the slots depends on the hardware on which 200 Series software is running. From the Configuration page, you can also manually configure information about cards on some platforms.

Slot Card Configuration

Use the **Card Configuration** page to view information about the cards installed in a switch. On some platforms, you can manually configure information about slots.

To access the page, click **System > Slot > Configuration** in the navigation menu.

Table 126 lists the fields that display when the slot contains a card.

Table 126: Slot Configuration Fields

Field	Description
Slot	Identifies the slot number.
Status	Whether the slot is empty or full.
Administrative State	Whether the slot is administratively enabled or disabled. For some devices, you can change the Administrative State when you add or edit slot information.
Power State	Whether the device is providing power to the slot. For some devices, you can change the Power State when you add or edit slot information.
Card Model	The model ID of the card configured for the slot.
Card Description	The description of the card configured for the slot.

Use the buttons to perform the following tasks:

- To preconfigure a card before adding it to a slot, click **Add** and configure the desired settings.
- To change slot or card settings, select the checkbox for the entry and click **Edit**.
- To delete a slot configuration entry from the list, select the checkbox for each entry to delete and click **Remove**.

Table 127: Card Configuration Fields

Field	Description
Unit	Indicates the unit in the stack for which data is to be displayed or configured.
Slot	Indicates the slot in the selected unit for which data is to be displayed or configured.
Card Description	The description of the card configured for the slot.
Card Index	Identifies the index number assigned to the card. This value is helpful when configuring the system by using SNMP .
Pluggable	If the value is True, the card can be administratively enabled or disabled. If the value is False, the Administrative State cannot be configured.
Power Down	If the value is True, the Power State can be administratively enabled or disabled. If the value is False, the Power State cannot be configured.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Click **Refresh** to redisplay the page with the current data from the switch.

Supported Cards

The **Supported Cards** page provides information about the cards that your platform supports.

To access this page, click **System > Slot > Supported Cards** in the navigation menu.

Table 128: Supported Card Fields

Field	Description
Card Index	Displays the index assigned to the selected card type.
Supported Cards	The menu contains the list of all cards that the system can support. To view information about a card, select it from the drop-down list. The screen refreshes, and the information about that card appears in the other fields on the page.
Card Type	Displays the hardware type of this supported card. This is a 32-bit data field.
Card Model ID	Displays the string to identify the model of the supported card.
Card Descriptor	Displays a data field used to identify the supported card.

Click **Refresh** to redisplay the most current information from the router.

Defining SNMP Parameters

[SNMP](#) provides a method for managing network devices. The device supports [SNMP version 1](#), [SNMP version 2](#), and [SNMP version 3](#).

SNMP v1 and v2

The *SNMP* agent maintains a list of variables, which are used to manage the device. The variables are defined in the *MIB (Management Information Base)*. The MIB presents the variables controlled by the agent. The SNMP agent defines the MIB specification format, as well as the format used to access the information over the network. Access rights to the SNMP agent are controlled by access strings.

SNMPv3

SNMPv3 also applies access control and a new traps mechanism to SNMPv1 and SNMPv2 PDUs. In addition, the *USM (User-based Security Model)* is defined for SNMPv3 and includes:

- **Authentication:** Provides data integrity and data origin authentication.
- **Privacy:** Protects against disclosure of message content. Cipher-Block-Chaining (CBC) is used for encryption. Either authentication is enabled on an SNMP message, or both authentication and privacy are enabled on an SNMP message. However privacy cannot be enabled without authentication.
- **Timeliness:** Protects against message delay or message redundancy. The SNMP agent compares incoming message to the message time information.
- **Key Management:** Defines key generation, key updates, and key use.

The device supports SNMP notification filters based on Object IDs (OID). OIDs are used by the system to manage device features. SNMP v3 supports the following features:

- Security
- Feature Access Control
- Traps

Authentication or Privacy Keys are modified in the SNMPv3 USM.

Use the SNMP page to define SNMP parameters. To display the SNMP page, click **System > Advanced Configuration > SNMP** in the navigation menu.

SNMP Community Configuration

Access rights are managed by defining communities on the SNMPv1, 2 Community page. When the community names are changed, access rights are also changed. *SNMP* Communities are defined only for SNMP v1 and SNMP v2.

Use the **SNMP Community Configuration** page to enable SNMP and Authentication notifications.

To access this page, click **System > Advanced Configuration > SNMP > Community** in the navigation menu.

Use the buttons to perform the following tasks:

- To add a community, click **Add** and configure the desired settings.
- To change information for an existing community, select the checkbox for the entry and click **Edit**.
- To delete a configured community from the list, select the checkbox for each entry to delete and click **Remove**.

Table 129: Community Configuration Fields

Field	Description
Community Name	<p>Contains the predefined and user-defined community strings that act as a password and are used to authenticate the SNMP management station to the device. A community string can contain a maximum of 20 characters. By default, the options available in the menu are as follows:</p> <ul style="list-style-type: none"> • public: This SNMP community has Read Only privileges and its status set to enable. • private: This SNMP community has Read/Write privileges and its status set to enable.
Community Access	Specifies the access control policy for the community.
Client IP Address	<p>Taken together, the Client IP Address and Client IP Mask denote a range of IP addresses from which SNMP clients may use that community to access this device. If either (IP Address or IP Mask) value is 0.0.0.0, access is allowed from any IP address. Otherwise, every client's IP address is ANDed with the mask, as is the Client IP Address, and, if the values are equal, access is allowed. For example, if the Client IP Address and Client IP Mask parameters are 192.168.1.0/255.255.255.0, then any client whose IP address is 192.168.1.0 through 192.168.1.255 (inclusive) will be allowed access. To allow access from only one station, use a Client IP Mask value of 255.255.255.255, and use that machine's IP address for Client IP Address.</p>
Client IP Mask	Along with the Client IP Address, the Client IP Mask denotes a range of IP addresses from which SNMP clients may use that community to access this device.
Access Mode	<p>Specify the access level for this community:</p> <ul style="list-style-type: none"> • Read-Only: The Community has read only access to the MIB objects configured in the view. • Read-Write: The Community has read/modify access to the MIB objects configured in the view.
Status	<p>Specify the status of this community:</p> <ul style="list-style-type: none"> • Enable: The community is enabled, and the Community Name must be unique among all valid Community Names or the set request will be rejected. • Disable: The Community is disabled and the Community Name becomes invalid.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Click **Remove** to delete the selected SNMP Community.

Trap Receiver v1/v2 Configuration

Use the **Trap Receiver v1/v2 Configuration** page to configure settings for each SNMPv1 or SNMPv2 management host that will receive notifications about traps generated by the device. The *SNMP* management host is also known as the SNMP trap receiver.

To access this page, click **System > Advanced Configuration > SNMP > Trap Receiver V1/V2** from the navigation menu.

Use the buttons to perform the following tasks:

- To add an SNMP trap receiver and configure its settings, click **Add** and complete the required information.
- To delete one or more SNMP trap receivers from the list, select each entry to delete and click **Remove**.

Table 130: Trap Receiver v1/v2 Configuration Fields

Field	Description
Host IP Address	The IP address of the SNMP management host that will receive traps generated by the device.
Community Name	The name of the SNMP community that includes the SNMP management host and the SNMP agent on the device.
Notify Type	The type of SNMP notification to send the SNMP management host: <ul style="list-style-type: none"> • Inform: An SNMP message that notifies the host when a certain event has occurred on the device. The message is acknowledged by the SNMP management host. This type of notification is not available for SNMPv1. • Trap: An SNMP message that notifies the host when a certain event has occurred on the device. The message is not acknowledged by the SNMP management host.
SNMP Version	The version of SNMP to use, which is either SNMPv1 or SNMPv2.
Timeout Value	The number of seconds to wait for an acknowledgment from the SNMP management host before resending an inform message.
Retries	The number of times to resend an inform message that is not acknowledged by the SNMP management host.
Filter	The name of the filter for the SNMP management host. The filter is configured by using the CLI and defines which <i>MIB</i> objects to include or exclude from the view. This field is optional.
UDP Port	The <i>UDP</i> port on the SNMP management host that will receive the SNMP notifications. If no value is specified when configuring a receiver, the default UDP port value is used.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Trap Receiver v3 Configuration

Use the **Trap Receiver V3 Configuration** page to configure settings for each SNMPv3 management host that will receive notifications about traps generated by the device. The *SNMP* management host is also known as the SNMP trap receiver.

To access this page, click **System > Advanced Configuration > SNMP > Trap Receiver V3** from the navigation menu.

Use the buttons to perform the following tasks:

- To add an SNMP trap receiver and configure its settings, click **Add** and complete the required information.

- To delete one or more SNMP trap receivers from the list, select each entry to delete and click **Remove**.

Table 131: Trap Receiver v3 Configuration Fields

Field	Description
Host IP Address	The IP address of the SNMP management host that will receive traps generated by the device.
User Name	The name of the SNMP user that is authorized to receive the SNMP notification.
Notify Type	The type of SNMP notification to send the SNMP management host: <ul style="list-style-type: none"> Inform: An SNMP message that notifies the host when a certain event has occurred on the device. The message is acknowledged by the SNMP management host. This type of notification is not available for SNMPv1. Trap: An SNMP message that notifies the host when a certain event has occurred on the device. The message is not acknowledged by the SNMP management host.
Security Level	The security level associated with the SNMP user, which is one of the following: <ul style="list-style-type: none"> No Auth No Priv: No authentication and no data encryption (no security). Auth No Priv: Authentication, but no data encryption. With this security level, users send SNMP messages that use an <i>MD5</i> key/password for authentication, but not a DES key/password for encryption. Auth Priv: Authentication and data encryption. With this security level, users send an MD5 key/password for authentication and a DES key/password for encryption.
Timeout Value	The number of seconds to wait for an acknowledgment from the SNMP management host before resending an inform message.
Retries	The number of times to resend an inform message that is not acknowledged by the SNMP management host.
Filter	The name of the filter for the SNMP management host. The filter is configured by using the CLI and defines which <i>MIB</i> objects to include or exclude from the view. This field is optional.
UDP Port	The <i>UDP</i> port on the SNMP management host that will receive the SNMP notifications. If no value is specified when configuring a receiver, the default UDP port value is used.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Supported MIBs

The **Supported MIBs** page lists each *MIB* that the system currently supports.

To access this page, click **System > Advanced Configuration > SNMP > Supported MIBs** in the navigation menu.

Table 132: Supported MIBs Fields

Field	Description
Name	The RFC number if applicable and the name of the MIB.
Description	The RFC title or MIB description.

Access Control Group

Use the **Access Control Group** page to configure SNMP access control groups. These SNMP groups allow network managers to assign different levels of authorization and access rights to specific device features and their attributes. The SNMP group can be referenced by the SNMP community to provide security and context for agents receiving requests and initiating traps as well as for management systems and their tasks. An SNMP agent will not respond to a request from a management system outside of its configured group, but an agent can be a member of multiple groups at the same time to allow communication with SNMP managers from different groups. Several default SNMP groups are preconfigured on the system.

To access this page, click **System > Advanced Configuration > SNMP > Access Control Group** in the navigation menu.

Use the buttons to perform the following tasks:

- To add an SNMP group, click **Add** and specify the desired setting.
- To remove one or more SNMP groups, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.

Table 133: Access Control Group Fields

Field	Description
Group Name	The name that identifies the SNMP group.
Context Name	The SNMP context associated with the SNMP group and its views. A user or a management application specifies the context name to get the performance information from the MIB objects associated with that context name. The Context EngineID identifies the SNMP entity that should process the request (the physical router), and the Context Name tells the agent in which context it should search for the objects requested by the user or the management application.
SNMP Version	The SNMP version associated with the group.
Security Level	The security level associated with the group, which is one of the following: <ul style="list-style-type: none"> • No Auth No Priv: No authentication and no data encryption (no security). This is the only Security Level available for SNMPv1 and SNMPv2 groups. • Auth No Priv: Authentication, but no data encryption. With this security level, users send SNMP messages that use an <u>MD5</u> key/password for authentication, but not a DES key/password for encryption. • Auth Priv: Authentication and data encryption. With this security level, users send an MD5 key/password for authentication and a DES key/password for encryption.

Table 133: Access Control Group Fields (continued)

Field	Description
Read	The level of read access rights for the group. The menu includes the available SNMP views. When adding a group, select the checkbox to allow the field to be configured, then select the desired view that restricts management access to viewing the contents of the agent.
Write	The level of write access rights for the group. The menu includes the available SNMP views. When adding a group, select the checkbox to allow the field to be configured, then select the desired view that permits management read-write access to the contents of the agent but not to the community.
Notify	The level of notify access rights for the group. The menu includes the available SNMP views. When adding a group, select the checkbox to allow the field to be configured, then select the desired view that permits sending SNMP traps or informs.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

User Security Model

The **User Security Model** page provides the capability to configure the SNMPv3 user accounts.

To access this page, click **System > Advanced Configuration > SNMP > User Security Model** in the navigation menu.

Use the buttons to perform the following tasks:

- To add a user, click **Add**. The **Add New SNMP User** dialog box opens. Specify the new account information in the available fields.
- To remove a user, select one or more table entries and click **Remove** to delete the selected entries.

Table 134: SNMP User Security Model Fields

Field	Description
Engine ID	Each SNMPv3 agent has an engine ID that uniquely identifies the agent in the device. If given this entry will be used only for packets whose engine ID is this. This field takes a hexadecimal string in the form 0102030405.
User Name	Specifies the name of the SNMP user being added for the User-based Security Model (USM). Each user name must be unique within the SNMP agent user list. A user name cannot contain any leading or embedded blanks.
Group Name	A SNMP group is a group to which hosts running the SNMP service belong. A group name parameter is simply the name of that group by which SNMP communities are identified. The use of a group name provides some security and context for agents receiving requests and initiating traps and does the same for management systems and their tasks. An SNMP agent won't respond to a request from a management system outside its configured group, but an agent can be a member of multiple groups at the same time. This allows for communications with SNMP managers from different groups.

Table 134: SNMP User Security Model Fields (continued)

Field	Description
Authentication Method	Specifies the authentication protocol to be used on authenticated messages on behalf of the specified user. <ul style="list-style-type: none"> • SHA: SHA protocol will be used. • MD5: MD5 protocol will be used. • None: No authentication will be used for this user.
Password	Specifies the password used to generate the key to be used in authenticating messages on behalf of this user. This parameter must be specified if the Authentication method parameter is not None.
Privacy	Specifies the privacy protocol to be used on encrypted messages on behalf of the specified user. This parameter is only valid if the Authentication method parameter is not NONE. <ul style="list-style-type: none"> • DES: DES protocol will be used. • None: No privacy protocol will be used.
Authentication Key	Specifies the password used to generate the key to be used in encrypting messages to and from this user. This parameter must be specified if the Privacy parameter is not None.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

SNMP View Entry

Use the **SNMP View Entry** page to configure *SNMP* views. These SNMP views allow network managers to control access to different parts of the MIB hierarchy permitting or denying access to objects. Once configured, views are associated to access control groups to complete access privileges.



Note

This feature is available for 220 switches only.

To access this page, click **System** > **Advanced Configuration** > **SNMP** > **View Entry** in the navigation menu.

Use the buttons to perform the following tasks:

- To add an SNMP view, click **Add**. Specify the desired settings and click **Submit**.
- To remove one or more SNMP views, select one or more views and click **Remove**. You must confirm the action before the views are removed.

Table 135: SNMP View Entry Fields

Field	Description
View Name	The name that identifies the SNMP view.
OID Tree	The ASN.1 subtree to be included or excluded from the view.
View Type	Type of access granted to the specified ASN.1 subtree: <ul style="list-style-type: none"> • Included: Access is granted to this subtree. • Excluded: Access is denied to this subtree.

Source Interface Configuration

Use **SNMP Trap Source Interface Configuration** page to specify the physical or logical interface to use as the *SNMP* client source interface. When an IP address is configured on the source interface, this address is used for all SNMP communications between the local SNMP client and the remote SNMP server. The IP address of the designated source interface is used in the IP header of SNMP management protocol packets. This allows security devices, such as firewalls, to identify all source packets coming from a specific device.

To access this page, click **System > Advanced Configuration > SNMP > Source Interface Configuration** in the navigation menu.

Table 136: SNMP Trap Source Interface Configuration Fields

Field	Description
Type	The type of interface to use as the source interface: <ul style="list-style-type: none"> • None: The primary IP address of the originating (outbound) interface is used as the source address. • Interface: The primary IP address of a physical port is used as the source address. • Loopback: The primary IP address of the loopback interface is used as the source address. A loopback is always reachable, as long as any routing interface is up. • VLAN: The primary IP address of a VLAN routing interface is used as the source address.
Interface	When the selected Type is Interface, select the physical port to use as the source interface.
Loopback Interface	When the selected Type is Loopback, select the loopback interface to use as the source interface.
VLAN ID	When the selected Type is VLAN, select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.
Tunnel ID	When the selected Type is Tunnel, select the tunnel interface to use as the source interface.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Viewing System Statistics

The pages in the Statistics folder contain a variety of information about the number and type of traffic transmitted from and received on the switch.

Switch Detailed Statistics

The **Switch Detailed Statistics** page shows detailed statistical information about the traffic the switch handles.

To access this page, click **System > Statistics > System > Switch** in the navigation menu.

Table 137: Switch Statistics Fields

Field	Description
Statistics	
Octets Without Error	The total number of octets (bytes) of data successfully transmitted or received by the processor (excluding framing bits but including FCS octets).
Packets Without Errors	The total number of packets including unicast, broadcast, and multicast packets, successfully transmitted or received by the processor.
Packets Discarded	The number of outbound (Transmit column) or inbound (Receive column) packets which were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher-layer protocol. A possible reason for discarding a packet could be to free up buffer space.
Unicast Packets	The number of subnetwork-unicast packets delivered to or received from a higher-layer protocol.
Multicast Packets	The total number of packets transmitted or received by the device that were directed to a multicast address. Note that this number does not include packets directed to the broadcast address.
Broadcast Packets	The total number of packets transmitted or received by the device that were directed to the broadcast address. Note that this number does not include multicast packets.
Status	
Current Usage	In the FDB Entries column, the value shows the number of learned and static entries in the MAC address table. In the VLANs column, the value shows the total number of static and dynamic VLANs that currently exist in the VLAN database.
Peak Usage	The highest number of entries that have existed in the MAC address table or VLAN database since the most recent reboot.
Maximum Allowed	The maximum number of statically configured or dynamically learned entries allowed in the MAC address table or VLAN database.
Static Entries	The current number of entries in the MAC address table or VLAN database that an administrator has statically configured.
Dynamic Entries	The current number of entries in the MAC address table or VLAN database that have been dynamically learned by the device.
Total Entries Deleted	The number of VLANs that have been created and then deleted since the last reboot. This field does not apply to the MAC address table entries.
System	

Table 137: Switch Statistics Fields (continued)

Field	Description
Interface	The interface index object value of the interface table entry associated with the Processor of this switch. This value is used to identify the interface when managing the device by using <i>SNMP</i> .
Time Since Counters Last Cleared	The amount of time in days, hours, minutes, and seconds, that has passed since the statistics for this device were last reset.

- Click **Refresh** to refresh the data on the screen with the present state of the data in the switch.
- Click **Clear Counters** to clear all the statistics counters, resetting all switch summary and detailed statistics to default values. The discarded packets count cannot be cleared.

Port Summary

The **Port Summary** page shows statistical information about the packets received and transmitted by each port and *LAG*.

To access this page, click **System > Statistics > System > Port Summary** in the navigation menu.

Table 138: Port Summary Fields

Field	Description
Interface	Identifies the port or LAG.
Rx Good	The total number of inbound packets received by the interface without errors.
Rx Errors	The number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol.
Rx Bcast	The total number of good packets received that were directed to the broadcast address. Note that this number does not include multicast packets.
Tx Good	The total number of outbound packets transmitted by the interface to its Ethernet segment without errors.
Tx Errors	The number of outbound packets that could not be transmitted because of errors.
Tx Collisions	The best estimate of the total number of collisions on this Ethernet segment.

- Click **Refresh** to refresh the data on the screen with the present state of the data in the switch.
- Click **Clear Counters** to clear all the statistics counters, resetting all summary and detailed statistics for this switch to default values. The discarded packets count cannot be cleared.
- Click **Clear All Counters** to clear counters for all switches in the stack.

Port Detailed Statistics

The **Port Detailed** page displays a variety of per-port traffic statistics.

To access this page, click **System > Statistics > System > Port Detailed** in the navigation menu.

Table 139: Port Detailed Statistics Fields

Field	Description
Interface	Use the drop-down menu to select the interface for which data is to be displayed or configured. For non-stacking systems, this field is Slot/Port.
Maximum Frame Size	The maximum Ethernet frame size the interface supports or is configured to support. The maximum frame size includes the Ethernet header, CRC, and payload.
Packet Lengths Received and Transmitted	
64 Octets	The total number of packets (including bad packets) received or transmitted that were 64 octets in length (excluding framing bits but including FCS octets).
65-127 Octets	The total number of packets (including bad packets) received or transmitted that were between 65 and 127 octets in length inclusive (excluding framing bits but including FCS octets).
128-255 Octets	The total number of packets (including bad packets) received or transmitted that were between 128 and 255 octets in length inclusive (excluding framing bits but including FCS octets).
256-511 Octets	The total number of packets (including bad packets) received or transmitted that were between 256 and 511 octets in length inclusive (excluding framing bits but including FCS octets).
512-1023 Octets	The total number of packets (including bad packets) received or transmitted that were between 512 and 1023 octets in length inclusive (excluding framing bits but including FCS octets).
1024-1518 Octets	The total number of packets (including bad packets) received or transmitted that were between 1024 and 1518 octets in length inclusive (excluding framing bits but including FCS octets).
1519-1522 Octets	The total number of packets (including bad packets) received or transmitted that were between 1519 and 1522 octets in length inclusive (excluding framing bits but including FCS octets).
1523-2047 Octets	The total number of packets (including bad packets) received or transmitted that were between 1523 and 2047 octets in length inclusive (excluding framing bits but including FCS octets).
2048-4095 Octets	The total number of packets (including bad packets) received or transmitted that were between 2048 and 4095 octets in length inclusive (excluding framing bits but including FCS octets).
4096-9216 Octets	The total number of packets (including bad packets) received or transmitted that were between 4096 and 9216 octets in length inclusive (excluding framing bits but including FCS octets).
Basic	
Unicast Packets	The Transmit column shows the total number of packets that higher-level protocols requested be transmitted to a subnetwork unicast address, including those that were discarded or not sent. The Receive column shows the number of subnetwork unicast packets delivered to a higher-layer protocol.
Multicast Packets	The Transmit column shows the total number of packets that higher-level protocols requested be transmitted to a multicast address, including those that were discarded or not sent. The Receive column shows the number of multicast packets delivered to a higher-layer protocol.

Table 139: Port Detailed Statistics Fields (continued)

Field	Description
Broadcast Packets	The Transmit column shows the total number of packets that higher-level protocols requested be transmitted to a broadcast address, including those that were discarded or not sent. The Receive column shows the number of broadcast packets delivered to a higher-layer protocol.
Total Packets (Octets)	The total number of octets of data (including those in bad packets) transmitted or received on the interface (excluding framing bits but including FCS octets). This object can be used as a reasonable estimate of Ethernet utilization. If greater precision is desired, the etherStatsPkts and etherStatsOctets objects should be sampled before and after a common interval.
Packets > 1518 Octets	The total number of packets transmitted or received by this interface that were longer than 1518 octets (excluding framing bits, but including FCS octets) and were otherwise well formed. This counter has a maximum increment rate of 815 counts per second at 10 Mb/s.
802.3x Pause Frames	The number of MAC Control frames transmitted or received by this interface with an opcode indicating the PAUSE operation. This counter does not increment when the interface is operating in half-duplex mode.
FCS Errors	The total number of packets transmitted or received by this interface that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, but had a bad Frame Check Sequence (FCS) with an integral number of octets.
Protocol	
STP BPDUs	The number of <i>STP BDU (Bridge Protocol Data Unit)</i> units transmitted or received by the interface.
RSTP BPDUs	The number of Rapid STP BPDUs transmitted or received by the interface.
MSTP BPDUs	The number of Multiple STP BPDUs transmitted or received by the interface.
SSTP BPDUs	The number of Shared STP BPDUs transmitted or received by the interface.
GVRP PDUs	The number of <i>GVRP (GARP VLAN Registration Protocol)</i> PDUs transmitted or received by the interface.
GMRP PDUs	The number of GARP Multicast Registration Protocol (GMRP) PDUs transmitted or received by the interface.
EAPOL Frames	The number of Extensible Authentication Protocol (EAP) over LAN (EAPOL) frames transmitted or received by the interface for IEEE 802.1X port-based network access control.
Advanced - Transmit	
Total Transmit Packets Discarded	The sum of single collision frames discarded, multiple collision frames discarded, and excessive frames discarded.
Single Collision Frames	A count of the number of successfully transmitted frames on a particular interface for which transmission is inhibited by exactly one collision.
Multiple Collision Frames	A count of the number of successfully transmitted frames on a particular interface for which transmission is inhibited by more than one collision.
Excessive Collision Frames	A count of frames for which transmission on a particular interface fails due to excessive collisions.

Table 139: Port Detailed Statistics Fields (continued)

Field	Description
Underrun Errors	The total number of frames discarded because the transmit FIFO buffer became empty during frame transmission.
GMRP Failed Registrations	The number of times attempted GMRP registrations could not be completed.
GVRP Failed Registrations	The number of times attempted GVRP registrations could not be completed.
Percent Utilization Transmitted (%)	The amount of link utilization, represented as a percentage of total link bandwidth, for the TX direction.
Advanced - Receive	
Total Packets Received Not Forwarded	The number of inbound packets which were chosen to be discarded to prevent them from being delivered to a higher-layer protocol, even though no errors had been detected. One possible reason for discarding such a packet is to free up buffer space.
Total Packets Received With MAC Errors	The total number of inbound packets that contained errors preventing them from being delivered to a higher-layer protocol.
Overruns	The total number of frames discarded as this port was overloaded with incoming packets, and could not keep up with the inflow.
Alignment Errors	The total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, but had a bad Frame Check Sequence (FCS) with a non-integral number of octets.
Jabbers Received	The total number of packets received that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error). Note that this definition of jabber is different than the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition where any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms.
Fragments Received	The total number of packets received that were less than 64 octets in length with ERROR CRC (excluding framing bits but including FCS octets).
Undersize Received	The total number of packets received that were less than 64 octets in length with GOOD CRC (excluding framing bits but including FCS octets).
Unacceptable Frame Type	The number of frames discarded from this interface due to being a frame type that the interface cannot accept.
Percent Utilization Received (%)	The amount of link utilization, represented as a percentage of total link bandwidth, for the RX direction.
Time Since Counters Last Cleared	The amount of time in days, hours, minutes, and seconds, that has passed since the statistics for this interface were last reset.

- Click **Clear Counters** to clear all the counters. This resets all statistics for this port to the default values.
- Click **Clear All Counters** to clear all the counters for all ports on the switch. The button resets all statistics for all ports to default values.
- Click **Refresh** to refresh the data on the screen and display the most current statistics.

Network DHCPv6

The **Network Port DHCPv6 Client Statistics** page displays the DHCPv6 client statistics values for the network interface. The DHCPv6 client on the device exchanges several different types of *UDP* messages with one or more network DHCPv6 servers during the process of acquiring address, prefix, or other relevant network configuration information from the server. The values indicate the various counts that have accumulated since they were last cleared.

To access this page, click **System > Statistics > System > Network DHCPv6** in the navigation menu.

Table 140: Network DHCPv6 Fields

Field	Description
Advertisement Packets Received	Number of DHCPv6 advertisement messages received from one or more DHCPv6 servers in response to the client's solicit message.
Reply Packets Received	Number of DHCPv6 reply messages received from one or more DHCPv6 servers in response to the client's request message.
Received Advertisement Packets Discarded	Number of DHCPv6 advertisement messages received from one or more DHCPv6 servers to which the client did not respond.
Received Reply Packets Discarded	Number of DHCPv6 reply messages received from one or more DHCPv6 servers to which the client did not respond.
Malformed Packets Received	Number of messages received from one or more DHCPv6 servers that were improperly formatted.
Total Packets Received	Total number of messages received from all DHCPv6 servers.
Solicit Packets Transmitted	Number of DHCPv6 solicit messages the client sent to begin the process of acquiring network information from a DHCPv6 server.
Request Packets Transmitted	Number of DHCPv6 request messages the client sent in response to a DHCPv6 server's advertisement message.
Renew Packets Transmitted	Number of renew messages the DHCPv6 client has sent to the server to request an extension of the lifetime of the information provided by the server. This message is sent to the DHCPv6 server that originally assigned the addresses and configuration information.
Rebind Packets Transmitted	Number of rebind messages the DHCPv6 client has sent to any available DHCPv6 server to request an extension of its addresses and an update to any other relevant information. This message is sent only if the client does not receive a response to the renew message.
Release Packets Transmitted	Number of release messages the DHCPv6 client has sent to the server to indicate that it no longer needs one or more of the assigned addresses.
Total Packets Transmitted	Total number of messages sent to all DHCPv6 servers.

Click **Clear Counters** to clear all the counters. This resets all statistics for this port to the default values.

Time Based Group Statistics

Use the **Time Based Group Statistics** page to define criteria for collecting time-based statistics for interface traffic. The time-based statistics can be useful for troubleshooting and diagnostics purposes. The statistics application uses the system clock for time-based reporting, so it is important to configure the system clock (manually or through *SNTP*) before using this feature.

To access this page, click **System > Statistics > Time Based > Group** in the navigation menu.

Use the buttons to perform the following tasks:

- To add a set of time-based traffic group statistics to collect, click **Add** and configure the desired settings.
- To delete one or more time-based statistics groups, select each entry to delete and click **Remove**.

Table 141: Time Based Group Statistics Fields

Field	Description
Group	The type of traffic statistics to collect for the group, which is one of the following: <ul style="list-style-type: none"> • Received: The number of packets received on the interfaces within the group. • Received Errors: The number of packets received with errors on the interfaces within the group. • Transmitted: The number of packets transmitted by the interfaces within the group. • Received Transmitted: The number of packets received and transmitted by the interfaces within the group. • Port Utilization: The percentage of total bandwidth used by the port within the specified time period. • Congestion: The percentage of time within the specified time range that the ports experienced congestion.
Time Range	The name of the periodic or absolute time range to use for data collection. The time range is configured by using the Time Range Summary and Time Range Entry Summary pages. The time range must be configured on the system before the time-based statistics can be collected.
Reporting Methods	The methods for reporting the collected statistics at the end of every configured time range interval. The available options are: <ul style="list-style-type: none"> • None: The statistics are not reported to the console or an external server. They can be viewed only by using the web interface or by issuing a CLI command. • Console: The statistics are displayed on the console. • E-Mail: The statistics are sent to an e-mail address. The SMTP server and e-mail address information is configured by using the appropriate Email Alerts pages. • Syslog: The statistics are sent to a remote syslog server. The syslog server information is configured on the Logging Hosts page.
Interfaces	The interface or interfaces on which data is collected. To select multiple interfaces when adding a new group, [Ctrl] + click each interface to include in the group.

- Click **Refresh** to refresh the data on the screen with the present state of the data in the switch.
- To add a set of time-based traffic group statistics to collect, click **Add** and configure the desired settings.
- To delete one or more time-based statistics groups, select each entry to delete and click **Remove**.

Time Based Flow Statistics

Use the **Time Based Flow Statistics** page to define criteria for collecting time-based statistics for specific traffic flows. The statistics include a per-interface hit count based on traffic that meets the match criteria configured in a rule for the interfaces included in the rule. The hit count statistics are collected only during the specified time range. The statistics application uses the system clock for time-based reporting. Configure the system clock (manually or through *SNTP*) before using the time-based statistics feature.

To access this page, click **System > Statistics > Time Based > Flow Based** in the navigation menu.

Use the buttons to perform the following tasks:

- To add a rule and define criteria for flow-based statistics that are collected within a time range, click **Add** and configure the desired settings.
- To delete one or more flow-based rules for time-based statistics, select each entry to delete and click **Remove**.
- Click **Refresh** to refresh the data on the screen with the present state of the data in the switch.
- To add a set of time-based traffic group statistics to collect, click **Add** and configure the desired settings.
- To delete one or more time-based statistics groups, select each entry to delete and click **Remove**.

Table 142: Time Based Flow Statistics Fields

Field	Description
Reporting Methods	The methods for reporting the collected statistics at the end of every configured interval. To change the reporting methods for all flow-based statistics rules, click the Edit icon and select one or more methods. To reset the field to its default value, click the Reset icon. The available reporting methods are: <ul style="list-style-type: none"> • None: The statistics are not reported to the console or an external server. They can be viewed only by using the web interface or by issuing a CLI command. • Console: The statistics are displayed on the console. • E-Mail: The statistics are sent to an e-mail address. The SNTP server and e-mail address information is configured by using the appropriate Email Alerts pages. • Syslog: The statistics are sent to a remote syslog server. The syslog server information is configured on the Logging Hosts page.
Rule Id	The number that identifies the flow-based statistics collection rule.
Time Range	The name of the periodic or absolute time range to use for data collection. The time range is configured by using the Time Range Summary and Time Range Entry Summary pages. The time range must be configured on the system before the time-based statistics can be collected.
Match Conditions	The criteria a packet must meet to match the rule.
Interfaces	The interface or interfaces on which the flow-based rule is applied. Only traffic on the specified interfaces is checked against the rule.

After you click **Add**, the **Time Based Flow Configuration** window opens and allows you to configure a rule for traffic flow statistics. The match conditions are optional, but the rule must specify at least one match condition. The following information describes the match criteria fields that are available in this window.

Table 143: Time Based Flow Configuration Fields

Field	Description
Match All	Select this option to indicate that all traffic matches the rule and is counted in the statistics. This option is exclusive to all other match criteria, so if Match All is selected, no other match criteria can be configured.
Source IP	The source IP address to match in the IPv4 packet header.
Destination IP	The destination IP address to match in the IPv4 packet header.
Source MAC	The source MAC address to match in the ingress frame header.
Destination MAC	The destination MAC address to match in the ingress frame header.
Source TCP Port	The TCP source port to match in the TCP header.
Destination TCP Port	The TCP destination port to match in the TCP header.
Source UDP Port	The UDP source port to match in the UDP header.
Destination UDP Port	The UDP destination port to match in the UDP header.

Time Based Statistics

Use the **Time Based Statistics** page to view time-based statistics collected for the configured traffic groups and flow-based rules.

To access this page, click **System > Statistics > Time Based > Statistics** in the navigation menu.

Table 144: Time Based Statistics Fields

Field	Description
ID	The traffic group name or flow-based rule ID associated with the rest of the statistics in the row.
Interface	The interface on which the statistics were reported.
Counter ID	For traffic group statistics, this field identifies the type of traffic.
Counter Value	For traffic group statistics, this field shows the number of packets of the type identified by the Counter Id field that were reported on the interface during the time range.
Port Utilization	For a port utilization traffic group, this field reports the percentage of the total available bandwidth used on the interface during the time range.
Hit Count	For flow-based statistics, this field reports the number of packets that matched the flow-based rule criteria during the time range.

Click **Refresh** to refresh the data on the screen with the present state of the data in the switch.

Using System Utilities

The System Utilities feature menu contains links to web pages that help you configure features that help you manage the switch.

System Reset

Use the **System Reset** page to reboot the system. If the platform supports stacking, you can reboot any or all of the switches in the stack from this page.

To access this page, click **System > Utilities > System Reset** in the navigation menu.

Table 145: System Reset Fields

Field	Description
Generate Core Dump before reset	Generates core dump file on demand.
Switch ID	Select the specific switch to reboot, or specify All to reboot all switches in the stack.
Reset (Button)	Initiates the system reboot action after displaying a confirmation message. Note that any configuration changes made since the last successful save are lost whenever a switch is rebooted. It is possible that the IP address of the switch will change. If this occurs, you will need to determine the new IP address to access the device using the web.

For Stacking platforms, you can select one or all switches in the stack to reboot from the drop-down menu. For platforms that do not support stacking, this field is not present.

Click **Reset** to initiate the system reboot. If you have not saved changes that you submitted since the last system reboot, the changes will not be applied to the system after the reboot.

Ping

Use the **Ping** page to tell the switch to send a ping request to a specified IP address. You can use this feature to check whether the switch can communicate with a particular network host.

To access this page, click **System > Utilities > Ping** in the navigation menu.

Table 146: Ping Fields

Field	Description
Hostname/IP Address	Enter the IP address or the host name of the station you want the switch to ping. The initial value is blank. This information is not retained across a power cycle.
Count	The number of <i>ICMP</i> echo request packets to send to the host.
Interval	The number of Seconds to wait between sending ping packets.
Size	The size of the ping packet, in bytes. Changing the size allows you to troubleshoot connectivity issues with a variety of packet sizes, such as large or very large packets.
Source	The source IP address or interface to use when sending the echo request packets. If source is not required, select None as source option.
IP Address	The source IP address to use when sending the Echo requests packets. This field is enabled when IP Address is selected as source option.
Interface	The interface to use when sending the Echo requests packets. This field is enabled when Interface is selected as source option.

Table 146: Ping Fields (continued)

Field	Description
Status	Displays the results of the ping.
Results	The results of the ping test, which includes information about the reply (if any) received from the host.
Start (Button)	Starts the ping test. The device sends the specified number of ping packets to the host.
Stop (Button)	Interrupts the current ping test.

Ping IPv6

Use the **Ping IPv6** page to tell the device to send one or more ping requests to a specified IPv6 host. You can use the ping request to check whether the device can communicate with a particular host on an IPv6 network. A ping request is an Internet Control Message Protocol version 6 (ICMPv6) echo request packet. The information you enter on this page is not saved as part of the device configuration.

To access this page, click **System > Utilities > Ping IPv6** in the navigation menu.

Table 147: Ping IPv6 Fields

Field	Description
Ping	Select either a global IPv6 address or a link local address to ping. A global address is routable over the Internet, while a link-local address is intended for communication only within the local network. Link local addresses have a prefix of fe80::/64.
Interface	This field displays only when Link Local is selected. Select an IPv6 interface to initiate the ping.
Host Name or IPv6 Address	Enter the global or link-local IPv6 address, or the DNS-resolvable host name of the station to ping. If the ping type is Link Local, you must enter a link-local address and cannot enter a host name.
Count	Enter the number of <i>ICMP</i> echo request packets to send to the host.
Interval	Enter the number of seconds to wait between sending ping packets.
Size	The size of the ping packet, in bytes. Changing the size allows you to troubleshoot connectivity issues with a variety of packet sizes, such as large or very large packets.
Source	The source IP address or interface to use when sending the echo request packets. If source is not required, select None as source option.
IPv6 Address	The source IPv6 address to use when sending the Echo requests packets. This field is enabled when IP Address is selected as source option.
Interface	The interface to use when sending the Echo requests packets. This field is enabled when Interface is selected as source option.
Results	The results of the ping test, which includes information about the reply (if any) received from the host.

Click **Submit** to send the specified number of pings. The results display in the **Ping Output** box.

TraceRoute

Use the **TraceRoute** page to determine the Layer 3 path a packet takes from the device to a specific IP address or hostname. When you initiate the TraceRoute command by clicking the Start button, the device sends a series of TraceRoute probes toward the destination. The results list the IP address of each Layer 3 device a probe passes through until it reaches its destination - or fails to reach its destination and is discarded. The information you enter on this page is not saved as part of the device configuration.

To access this page, click **System > Utilities > TraceRoute** in the navigation menu.

Table 148: TraceRoute Fields

Field	Description
Host Name or IP Address	The DNS-resolvable hostname or IP address of the system to attempt to reach.
Probes Per Hop	TraceRoute works by sending UDP packets with increasing Time-To-Live (TTL) values. Specify the number of probes sent with each TTL.
MaxTTL	The maximum Time-To-Live (TTL). The TraceRoute terminates after sending probes that can be Layer 3 forwarded this number of times. If the destination is further away, the TraceRoute will not reach it.
InitTTL	The initial Time-To-Live (TTL). This value controls the maximum number of Layer 3 hops that the first set of probes may travel.
MaxFail	The number of consecutive failures that terminate the TraceRoute. If the device fails to receive a response for this number of consecutive probes, the TraceRoute terminates.
Interval	The number of seconds to wait between sending probes.
Port	The <u>UDP</u> destination port number to be used in probe packets. The port number should be a port that the target host is not listening on, so that when the probe reaches the destination, it responds with an <u>ICMP</u> Port Unreachable message.
Size	The size of probe payload in bytes.
Source	Select None, IP Address, Interface, or Loopback as a source.
Status	The current status of the TraceRoute, which can be: <ul style="list-style-type: none"> • Not Started: The TraceRoute has not been initiated since viewing the page. • In Progress: The TraceRoute has been initiated and is running. • Stopped: The TraceRoute was interrupted by clicking the Stop button. • Done: The TraceRoute has completed, and information about the TraceRoute is displayed in the Results area.
Results	The results of the TraceRoute are displayed
Start (Button)	Initiates the TraceRoute.
Stop (Button)	Interrupts the running TraceRoute.

IP Address Conflict

Use the **IP Address Conflict** page to determine whether the IP address configured on the device is the same as the IP address of another device on the same LAN (or on the Internet, for a routable IP

address) and to help you resolve any existing conflicts. An IP address conflict can make both this system and the system with the same IP address unusable for network operation.

To access this page, click **System > Utilities > IP Address Conflict** in the navigation menu.

Table 149: IP Address Conflict Fields

Field	Description
IP Address Conflict Currently Exists	Whether a conflicting IP address has been detected since this status was last reset. <ul style="list-style-type: none"> False: No conflict detected (the subsequent fields on this page display as N/A). True: Conflict was detected (the subsequent fields on this page show the relevant information).
Last Conflicting IP Address	The device interface IP address that is in conflict. If multiple conflicts were detected, only the most recent occurrence is displayed.
Last Conflicting MAC Address	The MAC address of the remote host associated with the IP address that is in conflict. If multiple conflicts are detected, only the most recent occurrence is displayed.
Time Since Conflict Detected	The elapsed time (displayed in days, hours, minutes, and seconds) since the last address conflict was detected, provided the Clear History button has not yet been pressed.
Run Detection (Button)	Activates the IP address conflict detection operation in the system.
Clear History (Button)	Resets the IP address conflict detection status information that was last seen by the device.

Transfer

Use the **Transfer** page to upload files from the device to a remote system and to download files from a remote system to the device.

To access this page, click **System > Utilities > Transfer** in the navigation menu.

Table 150: Transfer Fields

Field	Description
Transfer Protocol	The protocol to use to transfer the file. Files can be transferred from the device to a remote system using TFTP, FTP, SCP or SFTP. Files can be transferred from a remote system to the device using HTTP, TFTP, FTP, SCP or SFTP.
Upload	To transfer a file from the device to a remote system using TFTP, FTP, SCP, or SFTP, click the upload icon in the same row as the desired transfer protocol. The File Upload window opens. Configure the information for the file transfer (described below), and click the upload icon to the right of the Progress field to begin the transfer.
Download	To transfer a file from a remote system to the device using HTTP, TFTP, FTP, SCP, or SFTP, click the download icon in the same row as the desired transfer protocol. The File Download window opens. Configure the information for the file transfer (described below), and click the download icon to the right of the Progress field to begin the transfer.

After you click the upload icon, the **File Upload** window opens.

The following information describes the fields in the **File Upload** window for all protocols.

Table 151: File Upload Fields

Field	Description
File Type	Specify the type of file to transfer from the device to a remote system. <ul style="list-style-type: none"> • Active Code: Select this option to transfer an active image. • Backup Code: Select this option to transfer a backup image. • Startup Configuration: Select this option to transfer a copy of the stored startup configuration from the device to a remote system. • Backup Configuration: Select this option to transfer a copy of the stored backup configuration from the device to a remote system. • Script File: Select this option to transfer a custom text configuration script from the device to a remote system. • CLI Banner: Select this option to transfer the file containing the text to be displayed on the CLI before the login prompt to a remote system. • Crash Log: Select this option to transfer the system crash log to a remote system. • Operational Log: Select this option to transfer the system operational log to a remote system. • Startup Log: Select this option to transfer the system startup log to a remote system. • Trap Log: Select this option to transfer the system trap records to a remote system. • Factory Defaults: Select this option to transfer the factory default configuration file to a remote system. • Error Log: Select this option to transfer the system error (persistent) log, which is also known as the event log, to a remote system. • Buffered Log: Select this option to transfer the system buffered (in-memory) log to a remote system.
Image	If the selected File Type is Code, specify whether to transfer the Active or Backup image to a remote system.
Server Address	Specify the IPv4 address, IPv6 address, or DNS-resolvable hostname of the remote server that will receive the file.
File Path	Specify the path on the server where you want to put the file.
File Name	Specify the name that the file will have on the remote server.
User Name	For FTP, SCP, and SFTP transfers, if the server requires authentication, specify the user name for remote login to the server that will receive the file.
Password	For FTP, SCP and SFTP transfers, if the server requires authentication, specify the password for remote login to the server that will receive the file.
Progress	Represents the completion percentage of the file transfer. The file transfer begins after you complete the required fields and click the upload icon to the right of this field.
Status	Provides information about the status of the file transfer.

After you click the download icon, the **File Download** window opens.

The following information describes the fields in the **File Download** window for all protocols.

Table 152: File Download Fields

Field	Description
File Type	<p>Specify the type of file to transfer to the device:</p> <ul style="list-style-type: none"> • Active Code: Select this option to transfer a new image to the device. The code file is stored as the active image. • Backup Code: Select this option to transfer a new image to the device. The code file is stored as the backup image. • Startup Configuration: Select this option to update the stored startup configuration file. If the file has errors, the update will be stopped. • Backup Configuration: Select this option to update the stored backup configuration file. If the file has errors, the update will be stopped. • Script File: Select this option to transfer a text-based configuration script to the device. You must use the command-line interface (CLI) to validate and activate the script. • CLI Banner: Select this option to transfer the CLI banner file to the device. This file contains the text to be displayed on the CLI before the login prompt. • IAS Users: Select this option to transfer an Internal Authentication Server (IAS) users database file to the device. The IAS user database stores a list of user name and (optional) password values for local port-based user authentication. • SSH-1 RSA Key File: Select this option to transfer an SSH-1 Rivest-Shamir-Adleman (RSA) key file to the device. SSH key files contain information to authenticate SSH sessions for remote CLI-based access to the device. • SSH-2 RSA Key PEM File: Select this option to transfer an SSH-2 Rivest-Shamir-Adleman (RSA) key file (PEM Encoded) to the device. • SSH-2 DSA Key PEM File: Select this option to transfer an SSH-2 Digital Signature Algorithm (DSA) key file (PEM Encoded) to the device. • Factory Defaults: Select this option to transfer the factory default configuration file to a remote system. • CA Root Certificate: Select this option to transfer an CA certificate file to the device. This will be used as the root certificate for one of the syslog servers. Based on the index number the file will be named accordingly. • Client Key: Select this option to transfer an client certificate file to the device. This will be used as the client certificate for one of the syslog servers. Based on the index number the file will be named accordingly. • Client SSL Certificate: Select this option to transfer an client key file to the device. Based on the index number the file will be named accordingly. • SSL Trusted Root Certificate PEM File: Select this option to transfer an SSL Trusted Root Certificate file (PEM Encoded) to the device. SSL files contain information to encrypt, authenticate, and validate HTTPS sessions. • SSL Server Certificate PEM File: Select this option to transfer an SSL Server Certificate file (PEM Encoded) to the device. • SSL DH Weak Encryption Parameter PEM File: Select this option to transfer an SSL Diffie-Hellman Weak Encryption Parameter file (PEM Encoded) to the device. • SSL DH Strong Encryption Parameter PEM File: Select this option to transfer an SSL Diffie-Hellman Strong Encryption Parameter file (PEM Encoded) to the device. • Public Key Image: Select this option to transfer the public key file used for code image validation to the device. • Public Key Config: Select this option to transfer the public key file used for configuration script validation to the device.

Table 152: File Download Fields (continued)

Field	Description
	To download SSH key files, SSH must be administratively disabled, and there can be no active SSH sessions. To download SSL related files, HTTPS must be administratively disabled.
Select File	If HTTP is the Transfer Protocol, browse to the directory where the file is located and select the file to transfer to the device. This field is not present if the Transfer Protocol is TFTP or FTP.
Certificate Index	Index used to name a related group of certificate (PEM) or key files.
Server Address	For TFTP, FTP, SCP, or SFTP transfers, specify the IPv4 address, IPv6 address, or DNS-resolvable hostname of the remote server.
File Path	For TFTP, FTP, SCP, or SFTP transfers, specify the path on the server where the file is located.
File Name	For TFTP, FTP, SCP, or SFTP transfers, specify the name of the file you want to transfer to the device.
User Name	For FTP, SCP, or SFTP transfers, if the server requires authentication, specify the user name for remote login to the server where the file resides.
Password	For FTP, SCP, or SFTP transfers, if the server requires authentication, specify the password for remote login to the server where the file resides.
Progress	Represents the completion percentage of the file transfer. The file transfer begins after you complete the required fields and click the download icon to the right of this field.
Digital Signature Verification	For Code and Startup Configuration file types, this option, when checked, will verify the file download with the digital signature.
Status	Provides information about the status of the file transfer.

Digital Signature Verification

Use the **Digital Signature Verification** page to configure digital signature verification on downloading files from a remote system to the device.

To access this page, click **System > Utilities > Digital Signature Verification** in the navigation menu.

Table 153: Digital Signature Verification Fields

Field	Description
Digital Signature Verification	Provides option to verify the digital signature of a downloaded file.
Code	Verify the digital signature of downloaded code image files.
Configuration	Verify the digital signature of downloaded configuration script files.

Core Dump

Use the **Core Dump** page to configure the Core Dump feature. A core dump contains the contents of the system's memory at the time of an abnormal termination, for use in debugging.

To access this page, click **System > Utilities > Core Dump** in the navigation menu.

Table 154: Core Dump Configuration Fields

Field	Description
Protocol	The protocol used to store the core dump file. User can select: <ul style="list-style-type: none"> • None: Disable Core Dump. • TFTP: Configure protocol to upload Core Dump to the TFTP server. • FTP: Configure protocol to upload Core Dump to the FTP server.
Core Dump File Name Prefix	Prefix for the Core Dump file name. If hostname is configured, it takes else while generating Core Dump file. The prefix length is 15 characters.
Use Host Name	To use hostname (or MAC if hostname is not configured) to name Core Dump file.
Use Time Stamp	To use timestamp to name Core Dump file.
TFTP IP Address	IP address of remote TFTP server to dump core file to external server.
FTP IP Address	IP address of remote FTP server to dump core file to external server.
FTP Username	Username of remote FTP server.
FTP Password	Password of remote FTP server.
File Path	File path to dump core file to TFTP server, NFS mount or USB device sub-directory.
Compression Mode	To enable or disable compression mode.
Switch Chip Registers Dump	To enable or disable switch-chip-register dump in case of an exception. The switch-chip-register dump is taken only for master unit and not for member units.
Stack IP Address Protocol	Protocol (<i>DHCP</i> or Static) to be used to configure service port when a unit has crashed. If configured as DHCP, the unit gets the IP address from DHCP server available in the network. If configured as Static, an IP address from the Core Dump Stack IP Address Pool is used.

Table 155: Core Dump Stack IP Address Pool Fields

Field	Description
IP Address	Static IP address to be assigned to individual unit's service port in the stack when the switch has crashed. This IP address is used to perform the core dump.
Host Mask	The subnet mask.
Default Router Address	The IP address of the router.

To add a stack IP address, click **Add** and configure an IP address, netmask, and gateway address.

To delete a configured stack IP, select each entry to delete, click **Remove**, and confirm the action.

Core Dump Test

Use the **Core Dump Test** page to test the core dump setup. For example, if protocol is configured as TFTP, it communicates with the TFTP server and informs the user if the TFTP server can be contacted.

To access this page, click **System > Utilities > Core Dump Test** in the navigation menu.

Table 156: Core Dump Test Fields

Field	Description
Status	Displays test status as <code>Ok</code> if test passes and <code>Error</code> if test fails.
Result	Displays detailed error information with logs.

4 Configuring Switching Information

Managing VLANs
Configuring UDLD
Private VLAN
Voice VLAN Configuration
Voice VLAN Interface
Port Auto Recovery
Creating MAC Filters
Configuring Dynamic ARP Inspection
GARP Configuration
Configuring DHCP Snooping
Configuring IPv6 DHCP Snooping
Configuring IGMP Snooping
Configuring IGMP Snooping Querier
Configuring MLD Snooping
Configuring MLD Snooping Querier
Creating Port Channels
Viewing Multicast Forwarding Database Information
Multicast VLAN Registration
Configuring Protected Ports
Configuring Spanning Tree Protocol
Mapping 802.1p Priority
Configuring Port Security
Managing LLDP
Loop Protection
Multiple Registration Protocol Configuration

Use the features in the Switching menu to define the switch's capabilities.

Managing VLANs

Adding Virtual LAN (VLAN) support to a Layer 2 switch offers some of the benefits of both bridging and routing. Like a bridge, a VLAN switch forwards traffic based on the Layer 2 header, which is fast, and like a router, it partitions the network into logical segments, which provides better administration, security and management of multicast traffic.

A VLAN is a set of end stations and the switch ports that connect them. You may have many reasons for the logical division, such as department or project membership. The only physical requirement is that the end station and the port to which it is connected both belong to the same VLAN.

Each VLAN in a network has an associated VLAN ID, which appears in the IEEE 802.1Q tag in the Layer 2 header of packets transmitted on a VLAN. An end station may omit the tag, or the VLAN portion of the tag, in which case the first switch port to receive the packet may either reject it or insert a tag using its default VLAN ID. A given port may handle traffic for more than one VLAN, but it can only support one default VLAN ID.

VLAN Status

Use the **VLAN Status** page to view information about the VLANs configured on your system.

To access this page, click **Switching > VLAN > Status** in the navigation menu.

Table 157: VLAN Status Fields

Field	Description
VLAN ID	The VLAN Identifier (VID) of the VLAN. The range of the VLAN ID is 1 to 3965.
VLAN Name	The name of the VLAN. VLAN ID 1 is always named Default.
VLAN Type	The VLAN type, which can be one of the following: <ul style="list-style-type: none"> • Default: (VLAN ID = 1) -- always present • Static: A VLAN you have configured • Dynamic: A VLAN created by GVRP registration that you have not converted to static, and that GVRP may therefore remove
RSPAN	List the status of RSPAN, enabled or disabled.

To add a VLAN, click **Add** and specify a VLAN ID (between 2 and 4093) in the available field.

To configure a name for a VLAN or to convert a dynamic VLAN to a static VLAN, select the entry to modify and click **Edit**. Then, configure the desired VLAN settings.

To remove one or more configured VLANs, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.

Click **Refresh** to display the latest information from the router.

VLAN Port Configuration

Use the **VLAN Port Configuration** page to configure a virtual LAN on a port.

To access this page, click **Switching > VLAN > Port Configuration** in the navigation menu.

Table 158: VLAN Port Configuration Fields

Field	Description
VLAN ID	The menu includes the VLAN ID for all VLANs configured on the device. To view or configure settings for a VLAN, be sure to select the correct VLAN from the menu.
Interface	Select the interface for which you want to display or configure data. Select All to set the parameters for all ports to same values.

Table 158: VLAN Port Configuration Fields (continued)

Field	Description
Status	The current participation mode of the interface in the selected VLAN. The value of the Status field differs from the value of the Participation field only when the Participation mode is set to Auto Detect. The Status is one of the following: <ul style="list-style-type: none"> • Include: The port is a member of the selected VLAN. • Exclude: The port is not a member of the selected VLAN.
Participation	The participation mode of the interface in the selected VLAN, which is one of the following: <ul style="list-style-type: none"> • Include: The port is always a member of the selected VLAN. This mode is equivalent to registration fixed in the IEEE 802.1Q standard. • Exclude: The port is never a member of the selected VLAN. This mode is equivalent to registration forbidden in the IEEE 802.1Q standard. • Auto Detect: The port can be dynamically registered in the selected VLAN through GVRP. The port will not participate in this VLAN unless it receives a GVRP request. This mode is equivalent to registration normal in the IEEE 802.1Q standard.
Tagging	The tagging behavior for all the ports in this VLAN, which is one of the following: <ul style="list-style-type: none"> • Tagged: The frames transmitted in this VLAN will include a VLAN ID tag in the Ethernet header. • Untagged: The frames transmitted in this VLAN will be untagged.

Use the buttons to perform the following tasks:

- To configure settings for one or more interfaces, select each interface to configure and click **Edit**. The same settings are applied to all selected interfaces.
- To apply the same settings to all interfaces, click **Edit All** and configure the desired settings.
- To reload the page and view the most current information, click **Refresh**.

VLAN Port Summary

Use the **VLAN Port Summary** page to view VLAN configuration information for all the ports on the system.

To access this page, click **Switching > VLAN > Port Summary** in the navigation menu.

Table 159: VLAN Port Summary Fields

Field	Description
Interface	Identifies the physical interface associated with the rest of the data in the row.
Port VLAN ID	The VLAN ID assigned to untagged or priority tagged frames received on this port. This value is also known as the Port VLAN ID (PVID). In a tagged frame, the VLAN is identified by the VLAN ID in the tag.

Table 159: VLAN Port Summary Fields (continued)

Field	Description
Acceptable Frame Types	<p>Indicates how the interface handles untagged and priority tagged frames. The options include the following:</p> <ul style="list-style-type: none"> • Admit All – Untagged and priority tagged frames received on the interface are accepted and assigned the value of the Port VLAN ID for this interface. • Only Tagged – The interface discards any untagged or priority tagged frames it receives. • Only Untagged – The interface discards any tagged frames it receives. <p>For all options, VLAN tagged frames are forwarded in accordance with the IEEE 802.1Q VLAN standard.</p>
Ingress Filtering	<p>Shows how the port handles tagged frames.</p> <ul style="list-style-type: none"> • Enable: A tagged frame is discarded if this port is not a member of the VLAN identified by the VLAN ID in the tag. • Disable: All tagged frames are accepted, which is the factory default.
Untagged VLANs	VLANs that are configured on the port to transmit egress packets as untagged.
Tagged VLANs	VLANs that are configured on the port to transmit egress packets as tagged.
Forbidden VLANs	When configuring port memberships in VLANs, you can specify one or more VLANs to be excluded from the available VLANs for the port. The forbidden VLANs list shows the VLANs to which the port cannot be assigned membership.
Dynamic VLANs	The list of VLANs of which the port became a member as result of the operations of dynamic VLAN protocols. When a VLAN is created as a dynamic VLAN, any port that is configured as switchport type Trunk or General automatically becomes a member of the VLAN, unless the VLAN port is excluded from the VLAN.
Priority	Identifies the default 802.1p priority assigned to untagged packets arriving at the port.

Use the buttons to perform the following tasks:

- To configure settings for one or more interfaces, select each interface to configure and click **Edit**. The same settings are applied to all selected interfaces.
- To apply the same settings to all interfaces, click **Edit All** and configure the desired settings.
- To reload the page and view the most current information, click **Refresh**.

Switchport Summary

Use the **Switchport Summary** page to configure switchport mode settings on interfaces. The switchport mode defines the purpose of the port based on the type of device it connects to and constraints the VLAN configuration of the port accordingly. Assigning the appropriate switchport mode helps simplify VLAN configuration and minimize errors.

To access this page, click **Switching > VLAN > Switchport Summary** in the navigation menu.

Table 160: VLAN Switchport Summary Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When editing information for one or more interfaces, this field identifies the interfaces that are being configured.
Switchport Mode	The switchport mode of the interface, which is one of the following: <ul style="list-style-type: none"> • Access: Access mode is suitable for ports connected to end stations or end users. Access ports participate only in one VLAN. They accept both tagged and untagged packets, but always transmit untagged packets. • Trunk: Trunk mode is intended for ports that are connected to other switches. Trunk ports can participate in multiple VLANs and accept both tagged and untagged packets. • General: General mode enables a custom configuration of a port. The user configures the General port VLAN attributes such as membership, PVID, tagging, ingress filter, etc., using the settings on the Port Configuration page. By default, all ports are initially configured in General mode.
Access VLAN ID	The access VLAN for the port, which is valid only when the port switchport mode is Access.
Native VLAN ID	The native VLAN for the port, which is valid only when the port switchport mode is Trunk.
Native VLAN Tagging	When enabled, if the trunk port receives untagged frames, it forwards them on the native VLAN with no VLAN tag. When disabled, if the port receives untagged frames, it includes the native VLAN ID in the VLAN tag when forwarding.
Trunk Allowed VLANs	The set of VLANs of which the port can be a member, when configured in Trunk mode. By default, this list contains all possible VLANs even if they have not yet been created.

Use the buttons to perform the following tasks:

- To configure settings for one or more interfaces, select each interface to configure and click **Edit**. The same settings are applied to all selected interfaces.
- To apply the same settings to all interfaces, click **Edit All** and configure the desired settings.
- To reload the page and view the most current information, click **Refresh**.

VLAN Internal Usage

Use the **VLAN Internal Usage** page to assign a Base VLAN ID for internal allocation of VLANs to the routing interface.

To access this page, click **Switching > VLAN > Internal Usage** in the navigation menu.

Table 161: VLAN Internal Usage Configuration Fields

Field	Description
Base VLAN ID	The first VLAN ID to be assigned to a port-based routing interface.
Allocation Policy	Determines whether VLAN IDs assigned to port-based routing interfaces start at the base and decrease in value (Descending) or start at the base and increase in value (Ascending).

Table 161: VLAN Internal Usage Configuration Fields (continued)

Field	Description
VLAN ID	The VLAN ID assigned to a port-based routing interface. The device automatically assigns an unused VLAN ID when the routing interface is created.
Routing Interface	The port-based routing interface associated with the VLAN.

If you change any information on the page, click **Submit** to apply the changes to the system.

Reset VLAN Configuration

Use the **Reset VLAN** page to return all VLAN parameters for all interfaces to the factory default values.

To access this page, click **Switching > VLAN > Reset** in the navigation menu.

When you click **Reset**, the screen refreshes, and you are asked to confirm the reset. Click **Reset** again to restore all default VLAN settings for the ports on the system.

RSPAN Configuration

Use the **RSPAN** page to configure the VLAN to use as the Remote Switched Port Analyzer (RSPAN) VLAN. RSPAN allows you to mirror traffic from multiple source ports (or from all ports that are members of a VLAN) from different network devices and send the mirrored traffic to a destination port (a probe port connected to a network analyzer) on a remote device. The mirrored traffic is tagged with the RSPAN VLAN ID and transmitted over trunk ports in the RSPAN VLAN.

To access this page, click **Switching > VLAN > RSPAN** in the navigation menu.

Table 162: RSPAN VLAN Configuration Fields

Field	Description
VLAN IDs	The VLANs configured on the system that are not currently enabled as Private VLANs. To enable a VLAN as a RSPAN VLAN, click the VLAN ID to select it (or [Ctrl] + click to select multiple VLAN IDs). Then, click the appropriate arrow to move the selected VLAN or VLANs to the RSPAN VLAN IDs window.
RSPAN VLAN IDs	The VLANs that are enabled as RSPAN VLAN. To disable a VLAN as a RSPAN VLAN, click the VLAN ID to select it (or [Ctrl] + click to select multiple VLAN IDs). Then, click the appropriate arrow to move the selected VLAN or VLANs to the VLAN IDs window.

Click **Refresh** to display the latest information from the router.

If you change any information on the page, click **Submit** to apply the changes to the system.

Configuring UDLD

The UDLD feature detects unidirectional links on physical ports by exchanging packets containing information about neighboring devices. The purpose of the UDLD feature is to detect and avoid

unidirectional links. A unidirectional link is a forwarding anomaly in a Layer 2 communication channel in which a bidirectional link stops passing traffic in one direction.

To access the **UDLD Configuration** page, click **Switching > UDLD > Configuration** in the navigation menu.

Table 163: UDLD Configuration Fields

Field	Description
Admin Mode	The administrative mode of UDLD on the device. UDLD must be administratively enabled on the device and on an interface for that interface to send UDLD messages. Additionally, UDLD must be enabled on the both sides of the link for the device to detect a unidirectional link.
Message Interval (Seconds)	The amount of time to wait between sending UDLD probe messages on ports that are in the advertisement phase.
Timeout Interval (Seconds)	The amount of time to wait to receive a UDLD message before considering the UDLD link to be unidirectional.

Click **Refresh** to display the latest information from the router.

If you change any information on the page, click **Submit** to apply the changes to the system.

UDLD Interface Configuration

Use the **UDLD Interface Configuration** page to configure the per-port UDLD settings.

To access this page, click **Switching > UDLD > Interface Configuration** in the navigation menu.

Use the buttons to perform the following tasks:

- To configure UDLD settings for one or more interfaces, select each interface to configure and click **Edit**. The same settings are applied to all selected interfaces.
- To reset all UDLD ports that have a UDLD Status of Shutdown, click **UDLD Port Reset**. If the global and interface UDLD administrative mode is enabled and the port link is up, the port restarts the exchange of UDLD messages with its link partner. The UDLD port status is Shutdown if UDLD has detected an unidirectional link and has put the port in a disabled state.

Table 164: UDLD Interface Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. In the Edit UDLD Interface Configuration window, this field identifies each interface that is being configured.
Admin Mode	The administrative mode of UDLD on the port.

Table 164: UDLD Interface Configuration Fields (continued)

Field	Description
UDLD Mode	<p>The UDLD mode for the port, which is one of the following:</p> <ul style="list-style-type: none"> • Normal: The state of the port is classified as Undetermined if an anomaly exists. An anomaly might be the absence of its own information in received UDLD messages or the failure to receive UDLD messages. An Undetermined state has no effect on the operation of the port. The port is not disabled and continues operating. When operating in UDLD normal mode, a port will be put into a disabled (Shutdown) state only in the following situations: <ul style="list-style-type: none"> • The UDLD PDU received from a partner does not have its own details (echo). • When there is a loopback, and information sent out on a port is received back exactly as it was sent. • Aggressive: The port is put into a disabled state for the same reasons that it occurs in normal mode. Additionally, a port in UDLD aggressive mode can be disabled if the port does not receive any UDLD echo packets even after bidirectional connection was established. If a bidirectional link is established, and packets suddenly stop coming from partner device, the UDLD aggressive-mode port assumes that link has become unidirectional.
UDLD Status	<p>The UDLD status on the port, which is one of the following:</p> <ul style="list-style-type: none"> • Not Applicable: The administrative status of UDLD is globally disabled or disabled on the interface. • Bidirectional: UDLD has detected a bidirectional link. • Shutdown: UDLD has detected a unidirectional link, and the port is in a disabled state. To clear the disabled state, click UDLD Port Reset. • Undetermined: UDLD has not collected enough information to determine the state of the port. • Unknown: The port link has physically gone down, but it is not because it was put in a disabled state by the UDLD feature.

Click **Refresh** to display the latest information from the router.

If you change any information on the page, click **Submit** to apply the changes to the system.

Private VLAN

Use this screen to add Virtual Local Area Networks (VLANs) to the device and to configure existing VLANs as private VLANs. Private VLANs provide Layer 2 isolation between ports that share the same broadcast domain. In other words, a private VLAN allows a VLAN broadcast domain to be partitioned into smaller point-to-multipoint subdomains. The ports participating in a private VLAN can be located anywhere in the Layer 2 network. Each subdomain is defined (represented) by a primary VLAN and a secondary VLAN. The primary VLAN ID is the same for all subdomains that belong to a private VLAN. The secondary VLAN ID differentiates subdomains from each another and provides Layer 2 isolation between ports that are members of the same private VLAN.

Private VLAN Configuration

To access the **Private VLAN Configuration** page, click **Switching > Private VLAN > Configuration** in the navigation menu.

Use the buttons to perform the following tasks:

- To add a VLAN, click **Add VLAN** and specify the VLAN ID(s) in the available field.
- To configure a private VLAN, select the entry to modify and click **Edit**. Then, configure the desired private VLAN setting.



Note

Default VLAN and management VLAN cannot be configured as a private VLANs and hence are not displayed on this page.

Table 165: Private VLAN Configuration Fields

Field	Description
VLAN ID	Displays the VLAN ID for which Private VLAN type is being set.
Type	<p>Use the Private VLAN Type menu to select the type of private VLAN. The factory default is Unconfigured.</p> <ul style="list-style-type: none"> • Primary: A private VLAN that forwards the traffic from the promiscuous ports to isolated ports, community ports, and other promiscuous ports in the same private VLAN. Only one primary VLAN can be configured per private VLAN. All ports within a private VLAN share the same primary VLAN. • Isolated: A secondary VLAN that carries traffic from isolated ports to promiscuous ports. Only one isolated VLAN can be configured per private VLAN. • Community: A secondary VLAN that forwards traffic between ports that belong to the same community and to the promiscuous ports. Multiple community VLANs can be configured per private VLAN. • Unconfigured: The VLAN is not configured as a private VLAN.

Click **Refresh** to display the latest information from the router.

Private VLAN Association

Use the **Private VLAN Association** page to configure the association between the primary VLAN and secondary VLANs. Associating a secondary VLAN with a primary VLAN allows host ports in the secondary VLAN to communicate outside the private VLAN. To configure a primary VLAN association, select the entry to modify and click **Edit**.

To access this page, click **Switching > Private VLAN > Association** in the navigation menu.

Use the buttons to perform the following tasks:

- To configure a primary VLAN association, select each entry to modify and click **Edit**.



Note

Isolated VLANs and Community VLANs are collectively called Secondary VLANs.

Table 166: Private VLAN Association Fields

Field	Description
Primary VLAN	The VLAN ID of each VLAN configured as a primary VLAN.
Isolated VLAN	The VLAN ID of the isolated VLAN associated with the primary VLAN. If the field is blank, no isolated VLAN has been associated with the primary VLAN. An isolated VLAN is a secondary VLAN that carries traffic from isolated ports to promiscuous ports. Only one isolated VLAN can be configured per private VLAN.
Community VLAN	The VLAN ID of each community VLAN associated with the primary VLAN. If the field is blank, no community VLANs have been associated with the primary VLAN. A community VLAN is a secondary VLAN that forwards traffic between ports that belong to the same community and to the promiscuous ports. Multiple community VLANs can be configured per private VLAN.

After you click **Edit**, the Edit Private VLAN Association window opens and allows you to create associations with the selected primary VLAN. The following information describes the field in this window.

- **Secondary VLAN** – The isolated or community VLANs that can be associated with the primary VLAN. Secondary VLANs that are already associated with a primary VLAN do not appear in the list and cannot be associated with another primary VLAN. To select multiple secondary VLANs, Ctrl + click each VLAN to associate with the primary VLAN.

Click **Refresh** to display the latest information from the router.

Private VLAN Interface

The **Private VLAN Interface** page allows you to configure the primary and secondary VLAN IDs for the host association mode. It also allows you to configure the port mode for the ports and LAGs that belong to a private VLAN and to configure associations between interfaces and primary/secondary private VLANs.

To access this page, click **Switching > Private VLAN > Interface** in the navigation menu.

Use the buttons to perform the following tasks:

- To configure the port mode and private VLAN-to-interface associations, select the entry to modify and click **Edit**.
- To remove the association between an interface and the primary/secondary private VLANs that the interface belongs to when it operates in host mode, select each interface with the association to clear and click **Remove Host Association**. You must confirm the action before the host association for the entry is cleared.
- To remove the association between an interface and the primary/secondary private VLANs that the interface belongs to when it operates in promiscuous mode, select each interface with the association to clear and click **Remove Promiscuous Association**. You must confirm the action before the promiscuous association for the entry is cleared.

Table 167: Private VLAN Interface Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When editing interface settings, this field identifies the interface being configured.
Mode	The private VLAN mode of the interface, which is one of the following: <ul style="list-style-type: none"> • General: The interface is in general mode and is not a member of a private VLAN. • Promiscuous: The interface belongs to a primary VLAN and can communicate with all interfaces in the private VLAN, including other promiscuous ports, community ports, and isolated ports. • Host: The interface belongs to a secondary VLAN and, depending upon the type of secondary VLAN, can either communicate with other ports in the same community (if the secondary VLAN is a community VLAN) and with the promiscuous ports or is able to communicate only with the promiscuous ports (if the secondary VLAN is an isolated VLAN).
Host Primary VLAN	The primary private VLAN the port is a member of when it is configured to operate in Host mode.
Host Secondary VLAN	The secondary private VLAN the port is a member of when it is configured to operate in Host mode. The secondary private VLAN is either an isolated or community VLAN.
Promiscuous Primary VLAN	The primary private VLAN in which the port is a member when it is configured to operate in Promiscuous mode.
Promiscuous Secondary VLAN	The secondary private VLAN the port is a member of when it is configured to operate in Promiscuous mode. The secondary private VLAN is either an isolated or community VLAN.
Operational Private VLAN	The primary and secondary operational private VLANs for the interface. The VLANs that are operational depend on the configured mode for the interface and the private VLAN type.

Click **Refresh** to display the latest information from the router.

Voice VLAN Configuration

The voice VLAN feature enables switch ports to carry voice traffic with defined settings so that voice and data traffic are separated when coming onto the port. A voice VLAN ensures that the sound quality of an IP phone is safeguarded from deterioration when data traffic on the port is high.

The inherent isolation provided by VLANs ensures that inter-VLAN traffic is under management control and that network-attached clients cannot initiate a direct attack on voice components. A [*QoS \(Quality of Service\)*](#) protocol based on the IEEE 802.1P [*CoS \(Class of Service\)*](#) protocol uses classification and scheduling to send network traffic from the switch in a predictable manner. The system uses the source MAC of the traffic traveling through the port to identify the IP phone data flow.

Voice VLAN is enabled per-port basis. A port can participate only in one voice VLAN at a time. The Voice VLAN feature is disabled by default.

To display the Voice VLAN Configuration page, click **Switching > Voice VLAN > Configuration**.

Table 168: Voice VLAN Configuration Fields

Field	Description
Voice VLAN Admin Mode	Click Enable or Disable to administratively turn the Voice VLAN feature on or off for all ports. The administrative mode of the Voice VLAN feature. When Voice VLAN is enabled globally and configured on interfaces that carry voice traffic, this feature can help ensure that the sound quality of an IP phone does not deteriorate when data traffic on the port is high.

- If you make any changes, click **Submit** to apply the change to the system.
- Click **Refresh** to display the latest information from the router.

Voice VLAN Interface

Use the **Voice VLAN Interface** page to configure the per-port settings for the Voice VLAN feature. When Voice VLAN is configured on a port that receives both voice and data traffic, it can help ensure that the voice traffic has priority.

Use the buttons to perform the following tasks:

- To configure Voice VLAN settings on a port, click **Add**. Select the interface to configure from the Interface menu, and then configure the desired settings.
- To change the Voice VLAN settings, select the interface to modify and click **Edit**.
- To remove the Voice VLAN configuration from one or more ports, select each entry to delete and click **Remove**.

To display this page, click **Switching > Voice VLAN > Interface Summary**.

Table 169: Voice VLAN Interface Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When adding a Voice VLAN configuration to a port, the Interface menu allows you to select the port to configure. Only interfaces that have not been configured with Voice VLAN settings can be selected from the menu.
Operational State	The operational status of the Voice VLAN feature on the interface. To be enabled, Voice VLAN must be globally enabled and enabled on the interface. Additionally, the interface must be up and have a link.
CoS Override Mode	The Class of Service override mode: <ul style="list-style-type: none"> • Enabled: The port ignores the 802.1p priority value in the Ethernet frames it receives from connected devices. • Disabled: The port trusts the priority value in the received frame.

Table 169: Voice VLAN Interface Fields (continued)

Field	Description
Voice VLAN Interface Mode	Indicates how an IP phone connected to the port should send voice traffic: <ul style="list-style-type: none"> • LAN ID: Forward voice traffic in the specified voice VLAN. • Dot1p: Tag voice traffic with the specified 802.1p priority value. • None: Use the settings configured on the IP phone to send untagged voice traffic. • Untagged: Send untagged voice traffic. • Disable: Operationally disables the Voice VLAN feature on the interface.
Voice VLAN Interface Value	When adding or editing Voice VLAN settings for an interface and either VLAN ID or Dot1p is selected as the Voice VLAN Interface Mode, specify the voice VLAN ID or the Dot1p priority value that the connected IP phone should use for voice traffic.

- If you make any changes, click **Submit** to apply the change to the system.
- Click **Refresh** to display the latest information from the router.

Port Auto Recovery

The Auto Recovery feature can automatically enable a disabled interface when the error conditions that caused the interface to be disabled are no longer detected. If Auto Recovery is not used (disabled), the interface remains disabled until an administrator manually enables it.

The switch supports an interface error disable feature that allows an interface to be automatically placed into a diagnostically disabled state when certain error conditions are detected on that interface. When an interface has been placed in a diagnostically disabled state, the interface is shut down, and no traffic is sent or received on that interface until it is either manually enabled by the administrator or re-enabled by the Auto Recovery feature after the recovery time interval has expired.

If the interface continues to encounter errors, it may be placed back into the diagnostically disabled state, and the interface will be disabled (link down). An interface in the diagnostically disabled state may also be manually recovered by enabling it from the Port Status page

Port Auto Recovery Configuration

Use the **Port Auto Recovery Configuration** page to allow a port to attempt to become re-enabled if it has been placed into a diagnostically disabled state due to the detection of certain error conditions.

To access this page, click **Switching > Auto Recovery > Configuration** in the navigation menu.

Table 170: Port Auto Recovery Configuration Fields

Field	Description
Auto Recovery Components	<p>This field lists all the components that support the Auto Recovery feature. For each component, you can enable or disable Auto Recovery.</p> <p>An interface in the diagnostic disabled state for the configured components is recovered (link up) when the recovery interval expires. If the interface continues to encounter errors (from any listed components), it may be placed back in the diagnostic disabled state, and the interface will be disabled (link down).</p> <p>Interfaces in the diagnostic disabled state may also be manually recovered by enabling them from the Port Summary page.</p> <p>Auto Recovery is available for the following components:</p> <ul style="list-style-type: none"> • ARP Inspection • BPDU Guard • BPDU Rate Limit • Broadcast Storm Control • Denial Of Service • <i>DHCP (Dynamic Host Configuration Protocol) Rate Limit</i> • Keepalive • MAC Locking • Multicast Storm Control • UDLD • Unicast Storm Control
Recovery Time	The auto recovery time interval. The auto recovery time interval is common for all components. The default value of the timer is 300 seconds and the range is from 30 to 86400.
D-Disabled Interface Status	This table displays the list of interfaces that are error disabled.
Interface	The interface which is error disabled.
Admin Mode	The administrative mode of the interface.
Port Status	Whether the link is up or down. The link is the physical connection between the port or trunk and the interface on another device.

Table 170: Port Auto Recovery Configuration Fields (continued)

Field	Description
Error Disable Reason	<p>If the device detects an error condition for an interface, then the device puts the interface in error disabled state by placing the interface in diagnostic disabled state. The interface can go into error disable state for one of the following reasons:</p> <ul style="list-style-type: none"> • ARP Inspection • BPDU Guard • BPDU Storm • Broadcast Storm • Denial Of Service • DHCP Rate Limit • Keepalive • MAC Locking • Multicast Storm • UDLD • Unicast Storm
Auto Recovery Time Left	<p>When Auto Recovery is enabled and the interface is placed in diagnostic disabled state, then a recovery timer starts for that interface. Once this timer expires, the device checks if the interface is in diagnostic disabled state. If yes, then the device enables the diagnostic disabled interface.</p>

Click **Submit** to apply the new configuration and cause the change to take effect. These changes will not be retained across a power cycle unless a Save configuration is performed.

Creating MAC Filters

Static MAC filtering allows you to associate a MAC address with a VLAN and set of source ports and destination ports. (The availability of source and destination port filters is subject to platform restrictions). Any packet with a static MAC address in a specific VLAN is admitted only if the ingress port is included in the set of source ports; otherwise the packet is dropped. If admitted, the packet is forwarded to all the ports in the destination list.

MAC Filter Configuration

Use the **MAC Filter Configuration** page to associate a MAC address with a VLAN and one or more source and/or destination ports

To access this page, click **Switching > Filters > MAC Filters** in the navigation menu.

Table 171: MAC Filter Configuration Fields

Field	Description
MAC Address	The MAC address of the filter. The destination MAC address of an Ethernet frame must match this value to be considered for the filter. When adding or editing a filter, note that you cannot configure the following MAC addresses in this field: <ul style="list-style-type: none"> • 00:00:00:00:00:00 • 01:80:C2:00:00:00 to 01:80:C2:00:00:0F • 01:80:C2:00:00:20 to 01:80:C2:00:00:21 • FF:FF:FF:FF:FF:FF
VLAN ID	The VLAN ID associated with the filter. The VLAN ID is used with the MAC address to fully identify the frames to filter.
Source Port Mask	The port(s) included in the inbound filter. If a frame with the MAC address and VLAN ID combination specified in the filter is received on a port in the Source Members list, it is forwarded to a port in the Destination Members list. If the frame that meets the filter criteria is received on a port that is not in the Source Members list, it is dropped. To add source ports to the filter, select one or more ports from the Available Port List field ([Ctrl] + click to select multiple ports). Then, use the appropriate arrow icon to move the selected ports to the Source Members field.
Destination Port Mask	The port(s) included in the outbound filter. A frame with the MAC address and VLAN ID combination specified in the filter is transmitted only out of ports in the list. To add destination ports to the filter, select one or more ports from the Available Port List field ([Ctrl] + click to select multiple ports). Then, use the appropriate arrow icon to add the selected ports to the Source Members field.

Adding MAC Filters

- 1 To add a MAC filter, click **Add** from the **MAC Filter** summary page.
- 2 Enter a valid MAC address and select a VLAN ID from the drop-down menu.
The VLAN ID drop-down menu only lists VLANs currently configured on the system.
- 3 Select one or more ports to include in the filter. Use **[Ctrl]** + click to select multiple ports.
- 4 Click **Submit** to apply the changes to the system.

Modifying MAC Filters

To change the port mask(s) for an existing filter, select the entry from the **MAC Filter** field, and click **Edit**. When you have completed the changes, click **Submit**.

To change the MAC address or VLAN associated with a filter, you must remove and re-create the filter.

Removing MAC Filters

To remove a filter, select it from the **MAC Filter** drop-down menu and click **Remove**.

Configuring Dynamic ARP Inspection

Dynamic ARP Inspection (DAI) is a security feature that rejects invalid and malicious ARP packets. DAI prevents a class of man-in-the-middle attacks, where an unfriendly station intercepts traffic for other

stations by poisoning the ARP caches of its unsuspecting neighbors. The miscreant sends ARP requests or responses mapping another station's IP address to its own MAC address.

DAI relies on *DHCP* snooping. DHCP snooping listens to DHCP message exchanges and builds a binding database of valid {MAC address, IP address, VLAN, and interface} tuples.

When DAI is enabled, the switch drops ARP packets whose sender MAC address and sender IP address do not match an entry in the DHCP snooping bindings database. You can optionally configure additional ARP packet validation.

DAI Configuration

Use the **DAI Configuration** page to configure global DAI settings.

To display this page, click **Switching > Dynamic ARP Inspection > Global** in the navigation menu.

Table 172: Dynamic ARP Inspection Global Configuration

Field	Description
Validate Source MAC	When this option is selected, DAI verifies that the sender hardware address in the ARP packet equals the source MAC address in the Ethernet header. If the addresses do not match, the ARP packet is dropped.
Validate Destination MAC	When this option is selected, DAI verifies that the target hardware address in the ARP packet equals the destination MAC address in the Ethernet header. If the addresses do not match, the ARP packet is dropped. This check applies only to ARP responses because the target MAC address is unspecified in ARP requests.
Validate IP	When this option is selected, DAI drops ARP packets with an invalid IP address. The following IP addresses are considered invalid: <ul style="list-style-type: none"> • 0.0.0.0 • 255.255.255.255 • All IP multicast addresses • All class E addresses (240.0.0.0/4) • Loopback addresses (in the range 127.0.0.0/8)

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

DAI VLAN Configuration

Use the **DAI VLAN Configuration** page to select the DAI-capable VLANs for which information is to be displayed or configured.

Use the buttons to perform the following tasks:

- To enable DAI on a VLAN and to configure the optional DAI settings, click **Add**.
- To change the DAI settings on VLAN, select the VLAN with the settings to update and click **Edit**.
- To disable DAI on one or more VLANs, select each entry to disable and click **Remove**. After confirming the action, the entries are removed from the table.

To display this page, click **Switching > Dynamic ARP Inspection > VLAN** in the navigation menu.

Table 173: Dynamic ARP Inspection VLAN Configuration

Field	Description
VLAN ID	Lists each VLAN that has been enabled for DAI. After you click Add , use the VLAN ID menu to select the VLAN on which to enable DAI. A VLAN does not need to exist on the system to be enabled for DAI.
Logging Invalid Packets	Whether DAI logging is enabled on this VLAN. When logging is enabled, DAI generates a log message whenever an invalid ARP packet is discovered and dropped.
ARP ACL Name	The name of the of ARP <i>ACL (Access Control List)</i> that the VLAN uses as the filter for ARP packet validation. The ARP ACL must already exist on the system to associate it with a DAI-enabled VLAN. ARP ACLs include permit rules only.
Static	Determines whether to use the <i>DHCP</i> snooping database for ARP packet validation if the packet does not match any ARP ACL rules. The options are as follows: <ul style="list-style-type: none"> • Enable: The ARP packet will be validated by the ARP ACL rules only. Packets that do not match any ARP ACL rules are dropped without consulting the DHCP snooping database. • Disable: The ARP packet needs further validation by using the entries in the DHCP Snooping database.

- Click **Submit** to apply the new configuration and cause the change to take effect. These changes will not be retained across a power cycle unless a Save configuration is performed.
- Click **Refresh** to refresh the page with the most current data from the switch.

DAI Interface Configuration

Use the **DAI Interface Configuration** page to select the DAI Interface for which information is to be displayed or configured.

To display this page, click **Switching > Dynamic ARP Inspection > Interface Configuration** in the navigation menu.

Table 174: Dynamic ARP Inspection Interface Configuration

Field	Description
Interface	The interface associated with the rest of the data in the row. In the Edit Interface Configuration window, this field identifies the interface that is being configured.
Trust State	Whether the DAI feature should check traffic on the interface for possible ARP packet violations. Trust state can be enabled or disabled after you select an interface and click Edit . This field has one of the following values: <ul style="list-style-type: none"> • Enabled: The interface is trusted. ARP packets arriving on this interface are forwarded without DAI validation. • Disabled: The interface is not trusted. ARP packets arriving on this interface are subjected to ARP inspection.
Rate Limit	The maximum rate for incoming ARP packets on the interface, in packets per second (pps). If the incoming rate exceeds the configured limit, the ARP packets are dropped. Rate limiting can be enabled or disabled after you select an interface and click Edit .

Table 174: Dynamic ARP Inspection Interface Configuration (continued)

Field	Description
Burst Interval	The number of consecutive seconds the interface is monitored for incoming ARP packet rate limit violations.
Rate Limiting	Select this option to allow the interface to drop ARP packets if the rate at which they are received on the interface exceeds the configured Rate Limit for the Burst Interval duration. If this option is clear, rate limiting is disabled.

- Click **Submit** to apply the new configuration and cause the change to take effect. These changes will not be retained across a power cycle unless a Save configuration is performed.
- Click **Refresh** to refresh the page with the most current data from the switch.

DAI ARP ACL Configuration

Use the [DAI ARP ACL Configuration](#) page to add or remove DAI ARP ACLs.

To display this Configuration page, click **Switching > Dynamic ARP Inspection > ACL Configuration** in the navigation menu.

Table 175: Dynamic ARP Inspection ARP ACL Configuration

Field	Description
ACL Name	The menu contains the ARP ACL names that exist on the system.
Sender IP Address	The IP address of a system that is permitted to send ARP packets. The ARP packet must match on both the Sender IP Address and Sender MAC Address values in the rule to be considered valid.
Sender MAC Address	The MAC address of a system that is permitted to send ARP packets. The ARP packet must match on both the Sender IP Address and Sender MAC Address values in the rule to be considered valid.
Remarks	One or more remarks configured for the selected ACL and associated with the rule during rule creation.

Use the buttons to perform the following tasks:

- To create an ARP ACL and configure the first rule, click **Add ACL**.
- To add a new rule to an existing ACL, click **Add Rule** and select the name of the ACL to update from the ACL Name menu. Then, configure the rule.
- To remove one or more ARP ACLs, select each entry to delete and click **Remove**.
- Click **Refresh** to refresh the page with the most current data from the switch.

DAI ARP ACL Rule Configuration

Use the [DAI ARP ACL Rule Configuration](#) page to add or remove DAI ARP ACL Rules.

To display this page, click **Add Rule** from the [Dynamic ARP Inspection ACL Configuration](#) page.

Table 176: Dynamic ARP Inspection ARP ACL Rule Configuration

Field	Description
Sender IP Address	To create a new rule for the selected ARP ACL, enter in this field the Sender IP Address match value for the ARP ACL.
Sender MAC Address	To create a new rule for the selected ARP ACL, enter in this field the Sender MAC Address match value for the ARP ACL.

Click **Submit** to add a new ARP ACL rule.

DAI ARP ACL Summary

Use the **DAI ARP ACL Configuration** page to configure ARP ACLs. An ARP ACL can contain one or more permit rules. Each rule contains the IP address and MAC address of a system allowed to send ARP packets. When an ARP ACL is associated with a DAI-enabled VLAN, and an ARP packet is received on an interface that is a member of that VLAN, DAI validates the address information in the ARP packet against the rules in the ACL. If the sender information in the ARP packet matches a rule in the ARP ACL, DAI considers the packet to be valid, and the packet is forwarded.

To display this page, click **Switching** > **Dynamic ARP Inspection** > **ACL Summary** in the navigation menu.

Table 177: Dynamic ARP Inspection ACL Summary Fields

Field	Description
ACL Name	The name of the ACL. Only the ACLs that appear in this column can be referenced by DNI-enabled VLANs.

Use the buttons to perform the following tasks:

- To add an ARP ACL, click **Add** and configure the ACL name.
- To configure rules for an ARP ACL, select the ACL to configure and click **Edit**. You are redirected to the Dynamic ARP Inspection ACL Configuration page for the selected ACL.
- To remove one or more ARP ACLs, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.
- Click **Refresh** to refresh the page with the most current data from the switch.

DAI Statistics

Use the **DAI Statistics** page to display the statistics per VLAN.

To display this page, click **Switching** > **Dynamic ARP Inspection** > **DAI Statistics** in the navigation menu.

Table 178: Dynamic ARP Inspection Statistics

Field	Description
VLAN ID	The DAI-enabled VLAN associated with the rest of the information in the row. When DAI is enabled on a VLAN, DAI is enabled on all interfaces that are members of that VLAN.
DHCP Drops	The number of ARP packets that have been dropped by DAI because no matching <i>DHCP</i> snooping binding entry was found in the DHCP snooping database.
ACL Drops	The number of ARP packets that have been dropped by DAI because the sender IP address and sender MAC address in the ARP packet did not match any rules in the ARP <i>ACL</i> associated with this VLAN. The static flag on this VLAN is enabled, which means ARP packets that fail to match an ARP ACL rule are dropped immediately and are not checked against the DHCP snooping database for further validation.
DHCP Permits	The number of ARP packets that were forwarded by DAI because a matching DHCP snooping binding entry was found in the DHCP snooping database.
ACL Permits	The number of ARP packets that were forwarded by DAI because the sender IP address and sender MAC address in the ARP packet matched a rule in the ARP ACL associated with this VLAN.
Bad Source MAC	The number of ARP packets that were dropped by DAI because the sender MAC address in ARP packet did not match the source MAC address in the Ethernet header.
Bad Dest MAC	The number of ARP packets that were dropped by DAI because the target MAC address in the ARP reply packet did not match the destination MAC address in the Ethernet header.
Invalid IP	The number of ARP packets that were dropped by DAI because the sender IP address in the ARP packet or target IP address in the ARP reply packet was invalid. The following IP addresses are considered invalid: <ul style="list-style-type: none"> • 0.0.0.0 • 255.255.255.255 • All IP multicast addresses • All class E addresses (240.0.0.0/4) • Loopback addresses (in the range 127.0.0.0/8)
Forwarded	The total number of valid ARP packets forwarded by DAI.
Dropped	The total number of invalid ARP packets dropped by DAI.

Click **Refresh** to refresh the page with the most current data from the switch.

GARP Configuration

Use this page to set the administrative mode for the features that use the Generic Attribute Registration Protocol (GARP), including GARP VLAN Registration Protocol (GVRP) and GARP Multicast Registration Protocol (GMRP). GARP is a general-purpose protocol that registers any network connectivity or membership-style information. GARP defines a set of switches interested in a given network attribute, such as VLAN ID or multicast address.

GARP Switch Configuration

To access the **GARP Switch Configuration** page, click **Switching > GARP > Switch** in the navigation menu.

Table 179: GARP Switch Configuration Fields

Field	Description
GVRP Mode	The administrative mode of GVRP on the system. When enabled, GVRP can help dynamically manage VLAN memberships on trunk ports.
GMRP Mode	The administrative mode of GMRP on the system. When enabled, GMRP can help control the flooding of multicast traffic by keeping track of group membership information. GMRP is similar to <i>IGMP (Internet Group Management Protocol)</i> snooping in its purpose, but IGMP snooping is more widely used. GMRP must be running on both the host and the switch to function properly.

Click **Refresh** to refresh the page with the most current data from the switch.

GARP Port Configuration

Use the **GARP Port Configuration** page to set the per-interface administrative mode for GARP VLAN Registration Protocol (GVRP) and GARP Multicast Registration Protocol (GMRP). On this page, you can also set the GARP timers for each interface. GVRP and GMRP use the same set of GARP timers to specify the amount of time to wait before transmitting various GARP messages.

To access this page, click **Switching > GARP > Port** in the navigation menu.

To change the GARP settings for one or more interfaces, select each interface to configure and click **Edit**. The same settings are applied to all selected interfaces.

Table 180: GARP Port Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When configuring one or more interfaces in the Edit GARP Port Configuration window, this field identifies the interfaces that are being configured.
GVRP Mode	The administrative mode of GVRP on the interface. When enabled, GVRP can help dynamically manage VLAN memberships on trunk ports. GVRP must also be enabled globally for the protocol to be active on the interface. When disabled, the protocol will not be active on the interface, and the GARP timers have no effect.
GMRP Mode	The administrative mode of GMRP on the interface. When enabled, GMRP can help control the flooding of multicast traffic by keeping track of group membership information. GMRP must also be enabled globally for the protocol to be active on the interface. When disabled, the protocol will not be active on the interface, and the GARP timers have no effect.
Join Timer (Centisecs)	The amount of time between the transmission of GARP PDUs registering (or re-registering) membership for a VLAN or multicast group.

Table 180: GARP Port Configuration Fields (continued)

Field	Description
Leave Timer (Centiseocs)	The amount of time to wait after receiving an unregister request for a VLAN or multicast group before deleting the associated entry. This timer allows time for another station to assert registration for the same attribute in order to maintain uninterrupted service.
Leave All Timer (Centiseocs)	The amount of time to wait before sending a LeaveAll PDU after the GARP application has been enabled on the interface or the last LeaveAll PDU was sent. A LeaveAll PDU indicates that all registrations will shortly be deregistered. Participants will need to rejoin in order to maintain registration.

Click **Refresh** to refresh the page with the most current data from the switch.

Configuring DHCP Snooping

DHCP snooping is a security feature that monitors DHCP messages between a DHCP client and DHCP servers to filter harmful DHCP messages and to build a bindings database of {MAC address, IP address, VLAN ID, port} tuples that are considered authorized. You can enable DHCP snooping globally and on specific VLANs, and configure ports within the VLAN to be trusted or untrusted. If a DHCP message arrives on an untrusted port, DHCP snooping filters messages that are not from authorized DHCP clients. DHCP server messages are forwarded only through trusted ports.

Global DHCP Snooping Configuration

Use the **Global DHCP Snooping Configuration** page to view and configure the global settings for *DHCP* Snooping.

To access this page, click **Switching > DHCP Snooping > Base > Global** in the navigation menu.

Table 181: Global DHCP Snooping Configuration Fields

Field	Description
DHCP Snooping Mode	The administrative mode of DHCP snooping on the device.
MAC Address Validation	Enables or disables the verification of the sender MAC address for DHCP snooping. When enabled, the device checks packets that are received on untrusted interfaces to verify that the MAC address and the DHCP client hardware address match. If the addresses do not match, the device drops the packet.

Click **Refresh** to refresh the page with the most current data from the switch.

DHCP Snooping VLAN Configuration

Use the **DHCP Snooping VLAN Configuration** page to view and configure the *DHCP* snooping settings on VLANs that exist on the device. DHCP snooping can be configured on switching VLANs and routing VLANs. For Layer 2 (non-routing) VLANs, DHCP snooping forwards valid DHCP client messages received on the VLANs. The message is forwarded on all trusted interfaces in the VLAN. When a DHCP packet is received on a routing VLAN, the DHCP snooping application applies its filtering rules and

updates the bindings database. If a client message passes filtering rules, the message is placed into the software forwarding path, where it may be processed by the DHCP relay agent, the local DHCP server, or forwarded as an IP packet.

To access this page, click **Switching > DHCP Snooping > Base > VLAN Configuration** in the navigation menu.

Use the buttons to perform the following tasks:

- To enable a VLAN for DHCP snooping, click **Add** and select the VLAN to administratively enable for DHCP snooping. To select multiple VLANs, **[Ctrl]** + click each VLAN to select.
- To disable DHCP snooping on one or more VLANs, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.

Table 182: DHCP Snooping VLAN Configuration Fields

Field	Description
VLAN ID	The VLAN ID that is enabled for DHCP snooping. In the Add DHCP Snooping VLAN Configuration window, this field lists the VLAN ID of all VLANs that exist on the device.
DHCP Snooping Mode	The current administration mode of DHCP snooping for the VLAN. Only VLANs that are enabled for DHCP snooping appear in the list.

Click **Refresh** to refresh the page with the most current data from the switch.

DHCP Snooping Interface Configuration

Use the **DHCP Snooping Interface Configuration** page to view and configure the *DHCP* snooping settings for each interface. The DHCP snooping feature processes incoming DHCP messages.

For DHCPRELEASE and DHCPDECLINE messages, the feature compares the receive interface and VLAN with the client's interface and VLAN in the binding database. If the interfaces do not match, the application logs the event (when logging of invalid packets is enabled) and drops the message. If MAC address validation is globally enabled, messages that pass the initial validation are checked to verify that the source MAC address and the DHCP client hardware address match. Where there is a mismatch, DHCP snooping logs the event (when logging of invalid packets is enabled) and drops the packet. To change the DHCP Snooping settings for one or more interfaces, select each entry to modify and click Edit. The same settings are applied to all selected interfaces.

To access this page, click **Switching > DHCP Snooping > Base > Interface Configuration** in the navigation menu.

Table 183: DHCP Snooping Interface Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When configuring the settings for one or more interfaces, this field identifies each interface that is being configured.
Trust State	The trust state configured on the interface. The trust state is one of the following: <ul style="list-style-type: none"> • Disabled: The interface is considered to be untrusted and could potentially be used to launch a network attack. DHCP server messages are checked against the bindings database. On untrusted ports, DHCP snooping enforces the following security rules: <ul style="list-style-type: none"> • DHCP packets from a DHCP server (DHCP OFFER, DHCPACK, DHCPNAK, DHCPRELEASEQUERY) are dropped. • DHCPRELEASE and DHCPDECLINE messages are dropped if the MAC address is in the snooping database but the binding's interface is other than the interface where the message was received. • DHCP packets are dropped when the source MAC address does not match the client hardware address if MAC Address Validation is globally enabled. • Enabled: The interface is considered to be trusted and forwards DHCP server messages without validation.
Log Invalid Packets	The administrative mode of invalid packet logging on the interface. When enabled, the DHCP snooping feature generates a log message when an invalid packet is received and dropped by the interface.
Rate Limit (pps)	The rate limit value for DHCP packets received on the interface. To prevent DHCP packets from being used as a DoS attack when DHCP snooping is enabled, the snooping application enforces a rate limit for DHCP packets received on untrusted interfaces. If the incoming rate of DHCP packets exceeds the value of this object during the amount of time specified for the burst interval, the port will be shutdown. You must administratively enable the port to allow it to resume traffic forwarding.
Burst Interval (Seconds)	The burst interval value for rate limiting on this interface. If the rate limit is unspecified, then burst interval has no meaning.

Click **Refresh** to refresh the page with the most current data from the switch.

DHCP Snooping Static Bindings

Use the **DHCP Snooping Static Bindings** page to view, add, and remove static bindings in the [DHCP](#) snooping bindings database.

To access this page, click **Switching > DHCP Snooping > Base > Static Bindings** in the navigation menu.

Table 184: DHCP Snooping Static Bindings Fields

Field	Description
Interface	The interface on which the DHCP client is authorized.
MAC Address	The MAC address associated with the DHCP client. This is the Key to the binding database.

Table 184: DHCP Snooping Static Bindings Fields (continued)

Field	Description
VLAN ID	The ID of the VLAN the client is authorized to use.
IP Address	The IP address of the client.

Click **Refresh** to refresh the page with the most current data from the switch.

DHCP Snooping Dynamic Bindings

Use the **DHCP Snooping Dynamic Bindings** page to view and clear dynamic bindings in the *DHCP* snooping bindings database. The DHCP snooping feature uses DHCP messages to build and maintain the bindings database. The bindings database includes data for clients only on untrusted ports. DHCP snooping creates a tentative binding from DHCP DISCOVER and REQUEST messages. Tentative bindings tie a client to an interface (the interface where the DHCP client message was received). Tentative bindings are completed when DHCP snooping learns the client's IP address from a DHCP ACK message on a trusted port. DHCP snooping removes bindings in response to DECLINE, RELEASE, and NACK messages. The DHCP snooping feature ignores the ACK messages as a reply to the DHCP Inform messages received on trusted ports.

To access this page, click **Switching > DHCP Snooping > Base > Dynamic Bindings** in the navigation menu.

Table 185: DHCP Snooping Dynamic Bindings Fields

Field	Description
Interface	The interface on which the DHCP client message was received.
MAC Address	The MAC address associated with the DHCP client that sent the message. This is the Key to the binding database.
VLAN ID	The VLAN ID of the client interface.
IP Address	The IP address assigned to the client by the DHCP server.
Lease Time	The remaining IP address lease time for the client.
Clear (Button)	To remove one or more entries in the database, select each entry to delete and click Clear . You must confirm the action before the entry is deleted.

Click **Refresh** to refresh the page with the most current data from the switch.

DHCP Snooping Persistent Configuration

Use the **DHCP Snooping Persistent Configuration** page to configure the persistent location of the *DHCP* snooping bindings database. The bindings database can be stored locally on the device or on a remote system somewhere else in the network. The device must be able to reach the IP address of the remote system to send bindings to a remote database.

To access this page, click **Switching > DHCP Snooping > Base > Persistent** in the navigation menu.

Table 186: DHCP Snooping Persistent Configuration Fields

Field	Description
Store	The location of the DHCP snooping bindings database, which is either locally on the device (Local) or on a remote system (Remote).
Remote IP Address	The IP address of the system on which the DHCP snooping bindings database will be stored. This field is available only if Remote is selected in the Store field.
Remote File Name	The file name of the DHCP snooping bindings database in which the bindings are stored. This field is available only if Remote is selected in the Store field.
Write Delay (Seconds)	The amount of time to wait between writing bindings information to persistent storage. This allows the device to collect as many entries as possible (new and removed) before writing them to the persistent file.

Click **Refresh** to refresh the page with the most current data from the switch.

DHCP Snooping Statistics

Use the **DHCP Snooping Statistics** page to view and clear per-interface statistics about the DHCP messages filtered by the *DHCP* snooping feature. Only interfaces that are enabled for DHCP snooping and are untrusted appear in the table.

To access this page, click **Switching > DHCP Snooping > Base > Statistics** in the navigation menu.

Table 187: DHCP Snooping Statistics Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.
MAC Verify Failures	The number of DHCP messages that were dropped because the source MAC address and client hardware address did not match. MAC address verification is performed only if it is globally enabled.
Client Ifc Mismatch	The number of packets that were dropped by DHCP snooping because the interface and VLAN on which the packet was received does not match the client's interface and VLAN information stored in the binding database.
DHCP Server Msgs Received	The number of DHCP server messages (DHCP OFFER, DHCP ACK, DHCP NAK, DHCP RELEASE QUERY) that have been dropped on an untrusted port.
Clear Counters (Button)	To reset the statistics to zero for one or more interfaces, select each interface with the data to reset and click Clear Counters. You must confirm the action before the entry is deleted.

Click **Refresh** to refresh the page with the most current data from the switch.

DHCP L2 Relay Global Configuration

Use the **DHCP L2 Relay Global Configuration** page to control the administrative mode of *DHCP* Layer 2 Relay on the device. In Layer 2 switched networks, there may be one or more infrastructure devices (for example, a switch) between the client and the Layer 3 Relay agent/DHCP server. In this instance, some of the client device information required by the Layer 3 Relay agent may not be visible to it. When this

happens, a Layer 2 Relay agent can be used to add the information that the Layer 3 Relay agent and DHCP server need to perform their roles in IP address configuration and assignment.

To access this page, click **Switching > DHCP Snooping > L2 Relay > Global** in the navigation menu.

Table 188: DHCP L2 Relay Global Configuration Fields

Field	Description
Admin Mode	The global mode of DHCP Layer 2 relay on the device. When enabled, the device can act as a DHCP Layer 2 relay agent. This functionality must also be enabled on each port you want this service to operate on.

Click **Refresh** to refresh the page with the most current data from the switch.

DHCP L2 Relay Interface Configuration

Use the **DHCP L2 Relay Interface Configuration** page to enable Layer 2 *DHCP* relay on individual ports. Note that L2 DHCP relay must also be enabled globally on the device. To change the DHCP Layer 2 relay settings for one or more interfaces, select each entry to modify and click Edit. The same settings are applied to all selected interfaces.

To access this page, click **Switching > DHCP Snooping > L2 Relay > Interface Configuration** in the navigation menu.

Table 189: DHCP L2 Relay Interface Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When configuring the settings for one or more interfaces, this field identifies each interface that is being configured.
L2 Relay Mode	The administrative mode of Layer 2 relay mode on the interface. When enabled, the interface can act as a DHCP relay agent and add information that the Layer 3 relay agent and DHCP server need to perform their roles in IP address configuration and assignment.
Trust Mode	The Layer 2 relay trust mode for the interface, which is one of the following: <ul style="list-style-type: none"> • Trusted: A trusted interface usually connects to other agents or servers participating in the DHCP interaction (for example, other Layer 2 or Layer 3 relay agents or servers). An interface in this mode always expects to receive DHCP packets that include Option 82 information. If Option 82 information is not included, these packets are discarded. • Untrusted: An untrusted interface is generally connected to clients. DHCP packets arriving on an untrusted interface are never expected to carry Option 82 and are discarded if they do.

Click **Refresh** to refresh the page with the most current data from the switch.

DHCP L2 Relay VLAN Configuration

Use the **DHCP L2 Relay VLAN Configuration** page to control the *DHCP* Layer 2 relay settings on a particular VLAN. The VLAN is identified by a service VLAN ID (S-VID), which a service provider uses to

identify a customer's traffic while traversing the provider network to multiple remote sites. The device uses the VLAN membership of the switch port client (the customer VLAN ID, or C-VID) to perform a lookup on a corresponding S-VID.

To access this page, click **Switching > DHCP Snooping > L2 Relay > VLAN Configuration** in the navigation menu.

Use the buttons to perform the following tasks:

- To enable a VLAN for DHCP Layer 2 relay, click **Add** and select the VLAN from the available menu.
- To update the DHCP Layer 2 relay settings for one or more VLANs, select each entry to update and click **Edit**. The same settings are applied to all selected VLANs.
- To disable one or more VLANs as DHCP Layer 2 relay agents, select the appropriate VLANs and click **Remove**. You must confirm the action.

Table 190: DHCP L2 Relay VLAN Configuration Fields

Field	Description
VLAN ID	The VLAN associated with the rest of the data in the row. When configuring the settings for one or more VLANs, this field identifies each VLAN that is being configured.
Circuit ID	The administrative mode of the circuit ID. When enabled, if a client sends a DHCP request to the device and the client is in a VLAN that corresponds to the S-VID, the device adds the client's interface number to the Circuit ID sub-option of Option 82 in the DHCP request packet. This enables the device to reduce the broadcast domain to which the server replies are switched when the broadcast bit is set for DHCP packets. When this bit is set, the server is required to echo Option 82 in replies. Since the circuit-id field contains the client interface number, the Layer 2 relay agent can forward the response to the requesting interface only, rather than to all ports in the VLAN).
Remote ID	The DHCP remote identifier string. When a string is entered here, if a client sends a DHCP request to the device and the client is in a VLAN that corresponds to the S-VID, the device adds the string to the Remote-ID sub-option of Option 82 in the DHCP request packet. This sub-option can be used by the server for parameter assignment. The content of this option is vendor-specific.

Click **Refresh** to refresh the page with the most current data from the switch.

DHCP L2 Relay Interface Statistics

The **DHCP L2 Relay Interface Statistics** page shows statistical information about the Layer 2 *DHCP* Relay requests received on trusted and untrusted interfaces. An interface is untrusted when the DHCP Layer 2 relay trust mode is disabled.

To access this page, click **Switching > DHCP Snooping > L2 Relay > Statistics** in the navigation menu.

Table 191: DHCP L2 Relay Interface Statistics Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.
Untrusted Server Messages With Option-82	The number of messages received on an untrusted interface from a DHCP server that contained Option 82 data. These messages are dropped.
Untrusted Client Messages With Option-82	The number of messages received on an untrusted interface from a DHCP client that contained Option 82 data. These messages are dropped.
Trusted Server Messages With Option-82	The number of messages received on a trusted interface from a DHCP server that contained Option 82 data. These messages are forwarded.
Untrusted Client Messages With Option-82	The number of messages received on a trusted interface from a DHCP client that contained Option 82 data. These messages are forwarded.
Clear (Button)	To reset the statistics to zero for one or more interfaces, select each interface with the data to reset and click Clear. You must confirm the action before the entry is deleted.

Click **Refresh** to refresh the page with the most current data from the switch.

Configuring IPv6 DHCP Snooping

IPv6 *DHCP* snooping is a security feature that monitors DHCPv6 messages between a DHCPv6 client and DHCPv6 servers to filter harmful DHCPv6 messages and to build a bindings database of {MAC address, IPv6 address, VLAN ID, port} tuples that are considered authorized. You can enable IPv6 DHCP snooping globally and on specific VLANs, and configure ports within the VLAN to be trusted or untrusted. If a DHCPv6 message arrives on an untrusted port, IPv6 DHCP snooping filters messages that are not from authorized DHCPv6 clients. DHCPv6 server messages are forwarded only through trusted ports.

Global IPv6 DHCP Snooping Configuration

Use the **Global IPv6 DHCP Snooping Configuration** page to view and configure the global settings for IPv6 DHCP Snooping.

To access this page, click **Switching > IPv6 DHCP Snooping > Base > Global** in the navigation menu.

Table 192: Global DHCP Snooping Configuration Fields

Field	Description
DHCP Snooping Mode	The administrative mode of IPv6 DHCP snooping on the device.
MAC Address Validation	Enables or disables the verification of the sender MAC address for IPv6 DHCP snooping. When enabled, the device checks packets that are received on untrusted interfaces to verify that the MAC address and the DHCPv6 client hardware address match. If the addresses do not match, the device drops the packet.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Click **Refresh** to refresh the page with the most current data from the switch.

IPv6 DHCP Snooping VLAN Configuration

Use the **IPv6 DHCP Snooping VLAN Configuration** page to view and configure the IPv6 *DHCP* snooping settings on VLANs that exist on the device. IPv6 DHCP snooping can be configured on switching VLANs and routing VLANs. For Layer 2 (non-routing) VLANs, IPv6 DHCP snooping forwards valid DHCPv6 client messages received on the VLANs. The message is forwarded on all trusted interfaces in the VLAN. When a DHCPv6 packet is received on a routing VLAN, the IPv6 DHCP snooping application applies its filtering rules and updates the bindings database. If a client message passes filtering rules, the message is placed into the software forwarding path, where it may be processed by the DHCPv6 relay agent, processed by the local DHCPv6 server, or forwarded as an IP packet.

To access this page, click **Switching > DHCP Snooping > Base > VLAN Configuration** in the navigation menu.

Use the buttons to perform the following tasks:

- To enable a VLAN for IPv6 DHCP snooping, click **Add** and select the VLAN to administratively enable for IPv6 DHCP snooping. To select multiple VLANs, **[Ctrl]** + click each VLAN to select.
- To disable IPv6 DHCP snooping on one or more VLANs, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.

Table 193: IPv6 DHCP Snooping VLAN Configuration Fields

Field	Description
VLAN ID	The VLAN ID that is enabled for IPv6 DHCP snooping. In the Add IPv6 DHCP Snooping VLAN Configuration window, this field lists the VLAN ID of all VLANs that exist on the device.
DHCP Snooping Mode	The current administration mode of IPv6 DHCP snooping for the VLAN. Only VLANs that are enabled for IPv6 DHCP snooping appear in the list.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Click **Refresh** to refresh the page with the most current data from the switch.

IPv6 DHCP Snooping Interface Configuration

Use the **IPv6 DHCP Snooping Interface Configuration** page to view and configure the IPv6 *DHCP* snooping settings for each interface. The IPv6 DHCP snooping feature processes incoming DHCPv6 messages.

For RELEASE and DECLINE messages, the feature compares the receive interface and VLAN with the client's interface and VLAN in the binding database. If the interfaces do not match, the application logs the event (when logging of invalid packets is enabled) and drops the message. If MAC address validation is globally enabled, messages that pass the initial validation are checked to verify that the source MAC address and the DHCPv6 client hardware address match. Where there is a mismatch, IPv6 DHCP snooping logs the event (when logging of invalid packets is enabled) and drops the packet. To

change the IPv6 DHCP Snooping settings for one or more interfaces, select each entry to modify and click **Edit**. The same settings are applied to all selected interfaces.

To access this page, click **Switching > IPv6 DHCP Snooping > Base > Interface Configuration** in the navigation menu.

Table 194: IPv6 DHCP Snooping Interface Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When configuring the settings for one or more interfaces, this field identifies each interface that is being configured.
Trust State	The trust state configured on the interface. The trust state is one of the following: <ul style="list-style-type: none"> • Disabled: The interface is considered to be untrusted and could potentially be used to launch a network attack. DHCPv6 server messages are checked against the bindings database. On untrusted ports, IPv6 DHCP snooping enforces the following security rules: <ul style="list-style-type: none"> • DHCPv6 packets from a DHCPv6 server (ADVERTISE, REPLY, and RECONFIGURE) are dropped. • RELEASE and DECLINE messages are dropped if the MAC address is in the snooping database but the binding's interface is other than the interface where the message was received. • DHCPv6 packets are dropped when the source MAC address does not match the client hardware address if MAC Address Validation is globally enabled. • Enabled: The interface is considered to be trusted and forwards DHCPv6 server messages without validation.
Log Invalid Packets	The administrative mode of invalid packet logging on the interface. When enabled, the IPv6 DHCP snooping feature generates a log message when an invalid packet is received and dropped by the interface.
Rate Limit (pps)	The rate limit value for DHCPv6 packets received on the interface. To prevent DHCPv6 packets from being used as a DoS attack when IPv6 DHCP snooping is enabled, the snooping application enforces a rate limit for DHCPv6 packets received on untrusted interfaces. If the incoming rate of DHCPv6 packets exceeds the value of this object during the amount of time specified for the burst interval, the port will be shut down. You must administratively enable the port to allow it to resume traffic forwarding.
Burst Interval (Seconds)	The burst interval value for rate limiting on this interface. If the rate limit is unspecified, then burst interval has no meaning.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Click **Refresh** to refresh the page with the most current data from the switch.

IPv6 DHCP Snooping Static Bindings

Use the [IPv6 DHCP Snooping Static Bindings](#) page to view, add, and remove static bindings in the IPv6 *DHCP* snooping bindings database.

To access this page, click **Switching > DHCP Snooping > Base > Static Bindings** in the navigation menu.

Use the buttons to perform the following tasks:

- To add a static entry to the IPv6 DHCP snooping bindings table, click **Add** and specify the desired settings.
- To remove one or more static entries, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.

Table 195: IPv6 DHCP Snooping Static Bindings Fields

Field	Description
Interface	The interface on which the DHCPv6 client is authorized.
MAC Address	The MAC address associated with the DHCPv6 client. This is the key to the binding database.
VLAN ID	The ID of the VLAN the client is authorized to use.
IP Address	The IPv6 address of the client.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Click **Refresh** to refresh the page with the most current data from the switch.

IPv6 DHCP Snooping Dynamic Bindings

Use the **IPv6 DHCP Snooping Dynamic Bindings** page to view and clear dynamic bindings in the IPv6 *DHCP* snooping bindings database. The IPv6 DHCP snooping feature uses DHCPv6 messages to build and maintain the bindings database. The bindings database includes data for clients only on untrusted ports.

IPv6 DHCP snooping creates a tentative binding from DHCPv6 SOLICIT and REQUEST messages. Tentative bindings tie a client to an interface (the interface where the DHCPv6 client message was received). Tentative bindings are completed when IPv6 DHCP snooping learns the client's IP address from a REPLY message on a trusted port. IPv6 DHCP snooping removes bindings in response to DECLINE and RELEASE messages.

To access this page, click **Switching > IPv6 DHCP Snooping > Base > Dynamic Bindings** in the navigation menu.

Table 196: IPv6 DHCP Snooping Dynamic Bindings Fields

Field	Description
Interface	The interface on which the DHCPv6 client message was received.
MAC Address	The MAC address associated with the DHCPv6 client that sent the message. This is the key to the binding database.
VLAN ID	The VLAN ID of the client interface.
IP Address	The IPv6 address assigned to the client by the DHCPv6 server.
Lease Time	The remaining IPv6 address lease time for the client.
Clear (Button)	To remove one or more entries in the database, select each entry to delete and click Clear . You must confirm the action before the entry is deleted.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Click **Refresh** to refresh the page with the most current data from the switch.

IPv6 DHCP Snooping Persistent Configuration

Use the **IPv6 DHCP Snooping Persistent Configuration** page to configure the persistent location of the IPv6 *DHCP* snooping bindings database. The bindings database can be stored locally on the device or on a remote system somewhere else in the network. The device must be able to reach the IP address of the remote system to send bindings to a remote database.

To access this page, click **Switching > IPv6 DHCP Snooping > Base > Persistent** in the navigation menu.

Table 197: IPv6 DHCP Snooping Persistent Configuration Fields

Field	Description
Store	The location of the IPv6 DHCP snooping bindings database, which is either locally on the device (Local) or on a remote system (Remote).
Remote IP Address	The IP address of the system on which the IPv6 DHCP snooping bindings database will be stored. This field is available only if Remote is selected in the Store field.
Remote File Name	The file name of the IPv6 DHCP snooping bindings database in which the bindings are stored. This field is available only if Remote is selected in the Store field.
Write Delay (Seconds)	The amount of time to wait between writing bindings information to persistent storage. This allows the device to collect as many entries as possible (new and removed) before writing them to the persistent file.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Click **Refresh** to refresh the page with the most current data from the switch.

DHCP Snooping Statistics

Use the **IPv6 DHCP Snooping Statistics** page to view and clear per-interface statistics about the DHCPv6 messages filtered by the IPv6 *DHCP* snooping feature. Only interfaces that are enabled for IPv6 DHCP snooping and are untrusted appear in the table.

To access this page, click **Switching > DHCP Snooping > Base > Statistics** in the navigation menu.

Table 198: IPv6 DHCP Snooping Statistics Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.
MAC Verify Failures	The number of DHCPv6 messages that were dropped because the source MAC address and client hardware address did not match. MAC address verification is performed only if it is globally enabled.

Table 198: IPv6 DHCP Snooping Statistics Fields (continued)

Field	Description
Client Ifc Mismatch	The number of packets that were dropped by IPv6 DHCP snooping because the interface and VLAN on which the packet was received does not match the client's interface and VLAN information stored in the binding database.
DHCP Server Msgs Received	The number of DHCPv6 server messages (ADVERTISE, REPLY, RECONFIGURE, RELAY-REPL) that have been dropped on an untrusted port.
Clear Counters (Button)	To reset the statistics to zero for one or more interfaces, select each interface with the data to reset and click Clear Counters . You must confirm the action before the entry is deleted.

Click **Refresh** to refresh the page with the most current data from the switch.

Configuring IGMP Snooping

IGMP snooping is a feature that allows a switch to forward multicast traffic intelligently on the switch. Multicast IP traffic is traffic that is destined to a host group. Host groups are identified by class D IP addresses, which range from 224.0.0.0 to 239.255.255.255. Based on the IGMP query and report messages, the switch forwards traffic only to the ports that request the multicast traffic. This prevents the switch from broadcasting the traffic to all ports and possibly affecting network performance.

A traditional Ethernet network may be separated into different network segments to prevent placing too many devices onto the same shared media. Bridges and switches connect these segments. When a packet with a broadcast or multicast destination address is received, the switch will forward a copy into each of the remaining network segments in accordance with the IEEE MAC Bridge standard. Eventually, the packet is made accessible to all nodes connected to the network.

This approach works well for broadcast packets that are intended to be seen or processed by all connected nodes. In the case of multicast packets, however, this approach could lead to less efficient use of network bandwidth, particularly when the packet is intended for only a small number of nodes. Packets will be flooded into network segments where no node has any interest in receiving the packet. While nodes will rarely incur any processing overhead to filter packets addressed to un-requested group addresses, they are unable to transmit new packets onto the shared media for the period of time that the multicast packet is flooded. The problem of wasting bandwidth is even worse when the LAN segment is not shared, for example in Full Duplex links.

Allowing switches to snoop IGMP packets is a creative effort to solve this problem. The switch uses the information in the IGMP packets as they are being forwarded throughout the network to determine which segments should receive packets directed to the group address.

Global Configuration and Status

Use the **IGMP Snooping Configuration and Status** page to enable *IGMP* snooping on the switch and view information about the current IGMP configuration.

To access this page, click **Switching > IGMP Snooping > Configuration** in the navigation menu.

Table 199: IGMP Snooping Global Configuration and Status Fields

Field	Description
Admin Mode	Select the administrative mode for IGMP Snooping for the switch from the drop-down menu. The default is disable.
Multicast Control Frame Count	Shows the number of multicast control frames that have been processed by the CPU.
Interfaces Enabled for IGMP Snooping	Lists the interfaces currently enabled for IGMP Snooping. To enable interfaces for IGMP snooping, see Interface Configuration on page 202.
Data Frames Forwarded by the CPU	Shows the number of data frames forwarded by the CPU.

Select **Enable** or **Disable** the **Admin Mode** field and click **Submit** to turn the feature on or off. Perform a save if you want the changes to remain in effect over a power cycle.

Interface Configuration

Use the [IGMP Snooping Interface Configuration](#) page to configure [IGMP](#) snooping settings on specific interfaces.

To access this page, click **Switching > IGMP Snooping > Interface Configuration** in the navigation menu.

Table 200: IGMP Snooping Interface Configuration Fields

Field	Description
Interface	Select the physical or LAG (Link Aggregation Group) interfaces to configure.
Admin Mode	Select the interface mode for the selected interface for IGMP Snooping for the switch from the drop-down menu. The default is disable.
Group Membership Interval	Specify the amount of time you want the switch to wait for a report for a particular group on a particular interface before it deletes that interface from the group. The valid range is 2 to 3600 seconds. The default is 260 seconds.
Max Response Time	Specify the amount of time you want the switch to wait after sending a query on an interface because it did not receive a report for a particular group on that interface. Enter a value greater or equal to 1 and less than the Group Membership Interval in seconds. The default is 10 seconds. The configured value must be less than the Group Membership Interval.
Multicast Router Present Expiration Time	Specify the amount of time you want the switch to wait to receive a query on an interface before removing it from the list of interfaces with multicast routers attached. Enter a value between 0 and 3600 seconds. The default is 0 seconds, indicating an infinite timeout or no expiration.
Fast Leave Admin Mode	Select the Fast Leave mode for the a particular interface from the drop-down menu. The default is Disable.

If you make any changes on the page, click **Submit** to apply the new settings to the switch.

VLAN Status

Use the **VLAN Status** page to enable or disable *IGMP* snooping on system VLANs and to view and configure per-VLAN IGMP snooping settings. Only VLANs that are enabled for IGMP snooping appear in the table.

To access this page, click **Switching > IGMP Snooping > VLAN Status** in the navigation menu.

Use the buttons to perform the following tasks:

- To enable IGMP snooping on a VLAN, click **Add** and configure the settings in the available fields.
- To change the IGMP snooping settings for an IGMP-snooping enabled VLAN, select the entry with the settings to change and click **Edit**.
- To disable IGMP snooping on one or more VLANs, select each VLAN to modify and click **Remove**. You must confirm the action before IGMP snooping is disabled on the selected VLANs. When IGMP snooping is disabled, the VLAN entry is removed from the table, but the VLAN itself still exists on the system.

Table 201: IGMP Snooping VLAN Status Fields

Field	Description
VLAN ID	The VLAN associated with the rest of the data in the row. When enabling IGMP snooping on a VLAN, use this menu to select the desired VLAN. Only VLANs that have been configured on the system and are not already enabled for IGMP snooping appear in the menu. When modifying IGMP snooping settings, this field identifies the VLAN that is being configured.
Admin Mode	The administrative mode of IGMP snooping on the VLAN. IGMP snooping must be enabled globally and on a VLAN for the VLAN to be able to snoop IGMP packets to determine which network segments should receive multicast packets directed to the group address.
Fast Leave Admin Mode	The administrative mode of Fast Leave on the VLAN. If Fast Leave is enabled, the VLAN can be immediately removed from the Layer 2 forwarding table entry upon receiving an IGMP leave message for a multicast group without first sending out MAC-based general queries.
Group Membership Interval (Seconds)	The number of seconds the VLAN should wait for a report for a particular group on the VLAN before the IGMP snooping feature deletes the VLAN from the group.
Max Response Time (Seconds)	The number of seconds the VLAN should wait after sending a query if does not receive a report for a particular group. The specified value should be less than the Group Membership Interval.
Multicast Router Expiration Time (Seconds)	The number of seconds the VLAN should wait to receive a query before it is removed from the list of VLANs with multicast routers attached.
Report Suppression Mode	The IGMPv1 and IGMPv2 report suppression mode. The device uses IGMP report suppression to limit the membership report traffic sent to multicast-capable routers. When this mode is enabled, the device does not send duplicate reports to the multicast router. Note that this mode is supported only when the multicast query has IGMPv1 and IGMPv2 reports. This feature is not supported when the query includes IGMPv3 reports. The options are as follows: <ul style="list-style-type: none"> • Enabled: Only the first IGMP report from all hosts for a group IGMP report is forwarded to the multicast routers. • Disabled: The device forwards all IGMP reports from all hosts in a multicast group to the multicast routers.

Click **Refresh** to refresh the page with the most current data from the switch.

Multicast Router Configuration

If a multicast router is attached to the switch, its existence can be learned dynamically. You can also statically configure a switch port as a multicast router interface. Use the **IGMP Snooping Multicast Router Configuration** page to manually configure an interface as a static multicast router interface.

To access this page, click **Switching > IGMP Snooping > Multicast Router Configuration** in the navigation menu.

Table 202: Multicast Router Configuration Fields

Field	Description
Interface	Select the physical or <i>LAG</i> interface to display.
Multicast Router	Set the multicast router status: <ul style="list-style-type: none"> • Enabled: The port is a multicast router interface. • Disabled: The port does not have a multicast router configured.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Multicast Router VLAN Status

If a multicast router is attached to the switch, its existence can be learned dynamically. You can also statically configure one or more VLANs on each interface to act as a multicast router interface, which is an interface that faces a multicast router or *IGMP* querier and receives multicast traffic.

To access the **Multicast Router VLAN Status** page, click **Switching > IGMP Snooping > Multicast Router VLAN Status** in the navigation menu.

Table 203: IGMP Snooping Multicast Router VLAN Status Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. Only interfaces that are configured with multicast router VLANs appear in the table.
VLAN IDs	The ID of the VLAN configured as enabled for multicast routing on the associated interface.

Use the buttons as follows:

- Click **Refresh** to refresh the page with the most current data from the switch.
- To disable all VLANs as multicast router interfaces for one or more physical ports or LAGs (link aggregation groups), select each entry to modify and click **Remove**.
- To enable or disable specific VLANs as multicast router interfaces for a physical port or LAG, use the **Add** and **Edit** buttons. A multicast router interface faces a multicast router or IGMP querier and receives multicast traffic.

Multicast Router VLAN Configuration

Use the **IGMP Snooping Multicast Router VLAN Configuration** page to enable or disable specific VLANs as multicast router interfaces for a physical port or LAG. A multicast router interface faces a multicast router or IGMP querier and receives multicast traffic.

To access this page, click **Switching > IGMP Snooping > Multicast Router VLAN Configuration** in the navigation menu.

Table 204: IGMP Snooping Multicast Router VLAN Configuration Fields

Field	Description
Interface	Select the port or LAG on which to enable or disable a VLAN multicast routing interface.
VLAN IDs	The VLANs configured on the system that are not currently enabled as multicast router interfaces on the selected port or LAG. To enable a VLAN as a multicast router interface, click the VLAN ID to select it (or [Ctrl] + click to select multiple VLAN IDs). Then, click the appropriate arrow to move the selected VLAN or VLANs to the Configured VLAN IDs window.
Configured VLAN IDs	The VLANs that are enabled as multicast router interfaces on the selected port or LAG. To disable a VLAN as a multicast router interface, click the VLAN ID to select it (or [Ctrl] + click to select multiple VLAN IDs). Then, click the appropriate arrow to move the selected VLAN or VLANs to the VLAN IDs window.

Click **Refresh** to refresh the page with the most current data from the switch.

Configuring IGMP Snooping Querier

Use this page to configure the global IGMP snooping querier settings on the device. IGMP snooping requires that one central switch or router periodically query all end-devices on the network to announce their multicast memberships. This central device is the IGMP querier. When Layer 3 IP multicast routing protocols are enabled in a network with IP multicast routing, the IP multicast router acts as the IGMP querier. However, if the IP-multicast traffic in a VLAN needs to be Layer 2 switched only, an IP-multicast router is not required. The IGMP snooping querier can perform the IGMP snooping functions on the VLAN.

IGMP Snooping Querier Configuration

To access the **IGMP Snooping Querier Configuration** page, click **Switching > IGMP Snooping Querier > Configuration** in the navigation menu.

Table 205: IGMP Snooping Querier Configuration Fields

Field	Description
Admin Mode	The administrative mode for the IGMP snooping querier on the device. When enabled, the IGMP snooping querier sends out periodic IGMP queries that trigger IGMP report messages from the switches that want to receive IP multicast traffic. The IGMP snooping feature listens to these IGMP reports to establish appropriate forwarding.
IP Address	The snooping querier address to be used as source address in periodic IGMP queries. This address is used when no IP address is configured on the VLAN on which the query is being sent.
IGMP Version	The IGMP protocol version used in periodic IGMP queries.
Query Interval (Seconds)	The amount of time the IGMP snooping querier on the device should wait between sending periodic IGMP queries.
Querier Expiry Interval (Seconds)	The amount of time the device remains in non-querier mode after it has discovered that there is a multicast querier in the network.

- If you make any changes to this page, click **Submit** to apply the changes to the system.
- Click **Refresh** to refresh the page with the most current data from the switch.

VLAN Configuration

Use the **IGMP Snooping Querier VLAN Configuration** page to enable the snooping querier feature on one or more VLANs and to configure per-VLAN *IGMP* snooping querier settings. Only VLANs that have the IGMP snooping querier feature enabled appear in the table.

To access this page, click **Switching > IGMP Snooping Querier > VLAN Configuration** in the navigation menu.

Use the buttons to perform the following tasks:

- To enable the IGMP snooping querier feature on a VLAN, click **Add** and specify the desired settings.
- To change the IGMP snooping querier settings for a VLAN, select the entry to modify and click **Edit**.
- To disable the IGMP snooping querier feature on one or more VLANs, select each entry to change and click **Remove**. You must confirm the action before the entry is deleted. Clicking this button does not remove the VLAN from the system.

Table 206: IGMP Snooping Querier VLAN Configuration Fields

Field	Description
VLAN ID	The VLAN on which the IGMP snooping querier is enabled. When enabling the IGMP snooping querier on a VLAN, use this menu to select the desired VLAN. Only VLANs that have been configured on the system and are not already enabled for the IGMP snooping querier appear in the menu. When modifying IGMP snooping querier settings, this field identifies the VLAN that is being configured.
Querier Election Participation	The participation mode for the IGMP snooping querier election process: <ul style="list-style-type: none"> • Enabled – The IGMP snooping querier on this VLAN participates in the querier election process when it discovers the presence of another querier in the VLAN. If the snooping querier finds that the other querier source IP address is lower than its own address, it stops sending periodic queries. If the snooping querier wins the election (because it has the lowest IP address), then it continues sending periodic queries. • Disabled – When the IGMP snooping querier on this VLAN sees other queriers of the same version in the VLAN, the snooping querier moves to the non-querier state and stops sending periodic queries.
Querier VLAN IP Address	The IGMP snooping querier address the VLAN uses as the source IP address in periodic IGMP queries sent on the VLAN. If this value is not configured, the VLAN uses the global IGMP snooping querier IP address.

Click **Refresh** to refresh the page with the most current data from the switch.

VLAN Status

Use the [IGMP Snooping Querier VLAN Status](#) page to view information about the *IGMP* snooping querier status for all VLANs that have the snooping querier enabled.

To access this page, click **Switching > IGMP Snooping Querier > VLAN Status** in the navigation menu.

Table 207: IGMP Snooping Querier VLAN Configuration Fields

Field	Description
VLAN ID	The VLAN associated with the rest of the data in the row. The table includes only VLANs that have the snooping querier enabled.
State	The operational state of the IGMP snooping querier on the VLAN, which is one of the following: <ul style="list-style-type: none"> • Querier: The snooping switch is the querier in the VLAN. The snooping switch will send out periodic queries with a time interval equal to the configured querier query interval. If the snooping switch sees a better querier (numerically lower) in the VLAN, it moves to non-querier mode. • Non-Querier: The snooping switch is in non-querier mode in the VLAN. If the querier expiry interval timer expires, the snooping switch moves into querier mode. • Disabled: The snooping querier is not operational on the VLAN. The snooping querier moves to the disabled mode when IGMP snooping is not operational on the VLAN, when the querier address is not configured, or the network management address is not configured.

Table 207: IGMP Snooping Querier VLAN Configuration Fields (continued)

Field	Description
Version	The operational IGMP protocol version of the querier.
Last IP Address	The IP address of the last querier from which a query was snooped on the VLAN.
Last Version	The IGMP protocol version of the last querier from which a query was snooped on the VLAN.
Max Response Time (Seconds)	The maximum response time to be used in the queries that are sent by the snooping querier.

Click **Refresh** to refresh the page with the most current data from the switch.

Configuring MLD Snooping

Use this page to enable Multicast Listener Discovery (MLD) snooping on the device and to view global status information. In IPv4, Layer 2 switches can use *IGMP* snooping to limit the flooding of multicast traffic by dynamically configuring Layer 2 interfaces so that multicast traffic is forwarded to only those interfaces associated with IP multicast address.

In IPv6 networks, MLD snooping performs a similar function. With MLD snooping, IPv6 multicast data is selectively forwarded to a list of ports intended to receive the data (instead of being flooded to all of the ports in a VLAN). This list is constructed by



Note

This feature is available for 220 switches only.

Global Configuration and Status

To access the **MLD Snooping Configuration and Status** page, click **Switching > MLD Snooping > Configuration** in the navigation menu.

Table 208: MLD Snooping Configuration and Status Fields

Field	Description
MLD Snooping Admin Mode	The administrative mode of MLD snooping on the device.
Multicast Control Frame Count	The number of multicast control frames that have been processed by the CPU.
Interfaces Enabled for MLD Snooping	One or more interfaces on which MLD snooping is administratively enabled. MLD snooping must be enabled globally and on an interface for the interface to be able to snoop MLD packets to determine which segments should receive multicast packets directed to the group address.
VLANs Enabled for MLD Snooping	One or more VLANs on which MLD snooping is administratively enabled.

Select **Enable** or **Disable** in the **MLD Snooping > Admin Mode** field and click **Submit** to turn the feature on or off. Perform a save if you want the changes to remain in effect over a power cycle.

Interface Configuration

Use the **MLD Snooping Interface Configuration** page to configure MLD snooping settings on specific interfaces. To configure the settings for one or more interfaces, select each entry to modify and click **Edit**. The same MLD snooping settings are applied to all selected interfaces.

To access this page, click **Switching > MLD Snooping > Interface Configuration** in the navigation menu.

Table 209: MLD Snooping Interface Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When configuring MLD snooping settings, this field identifies each interface that is being configured.
Admin Mode	The administrative mode of MLD snooping on the interface. MLD snooping must be enabled globally and on an interface for the interface to be able to snoop MLD packets to determine which segments should receive multicast packets directed to the group address.
Group Membership Interval	The number of seconds the interface should wait for a report for a particular group on the interface before the MLD snooping feature deletes the interface from the group.
Max Response Time	The number of seconds the interface should wait after sending a query if it does not receive a report for a particular group. The specified value should be less than the Group Membership Interval.
Multicast Router Present Expiration Time	The number of seconds the interface should wait to receive a query before it is removed from the list of interfaces with multicast routers attached.
Fast Leave Admin Mode	The administrative mode of Fast Leave on the interface. If Fast Leave is enabled, the interface can be immediately removed from the Layer 2 forwarding table entry upon receiving an MLD leave message for a multicast group without first sending out MAC-based general queries.

If you make any changes on the page, click **Submit** to apply the new settings to the switch.

VLAN Status

Use the **VLAN Status** page to enable or disable MLD snooping on system VLANs and to view and configure per-VLAN MLD snooping settings. Only VLANs that are enabled for MLD snooping appear in the table.

To access this page, click **Switching > MLD Snooping > VLAN Status** in the navigation menu.

Use the buttons to perform the following tasks:

- To enable MLD snooping on a VLAN, click **Add** and configure the settings in the available fields.
- To change the MLD snooping settings for an MLD-snooping enabled VLAN, select the entry with the settings to change and click **Edit**.
- To disable MLD snooping on one or more VLANs, select each VLAN to modify and click **Remove**. You must confirm the action before MLD snooping is disabled on the selected VLANs. When MLD snooping is disabled, the VLAN entry is removed from the table, but the VLAN itself still exists on the system.

Table 210: MLD Snooping VLAN Status Fields

Field	Description
VLAN ID	The VLAN associated with the rest of the data in the row. When enabling MLD snooping on a VLAN, use this menu to select the desired VLAN. Only VLANs that have been configured on the system and are not already enabled for MLD snooping appear in the menu. When modifying MLD snooping settings, this field identifies the VLAN that is being configured.
Admin Mode	The administrative mode of MLD snooping on the VLAN. MLD snooping must be enabled globally and on a VLAN for the VLAN to be able to snoop MLD packets and determine which network segments should receive multicast packets directed to the group address.
Group Membership Interval	The number of seconds the VLAN should wait for a report for a particular group on the VLAN before the MLD snooping feature deletes the VLAN from the group.
Max Response Time	The number of seconds the VLAN should wait after sending a query if does not receive a report for a particular group. The specified value should be less than the Group Membership Interval.
Multicast Router Expiration Time	The number of seconds the VLAN should wait to receive a query before it is removed from the list of VLANs with multicast routers attached.
Fast Leave Admin Mode	The administrative mode of Fast Leave on the VLAN. If Fast Leave is enabled, the VLAN can be immediately removed from the Layer 2 forwarding table entry upon receiving an MLD leave message for a multicast group without first sending out MAC-based general queries.

Click **Refresh** to refresh the page with the most current data from the switch.

Multicast Router Configuration

Use the **MLD Snooping Multicast Router Configuration** page to manually configure an interface as a static MLD snooping multicast router interface. If a multicast router is attached to the switch, its existence can be learned dynamically. You can also statically configure an interface as a multicast router interface, which is an interface that faces a multicast router or MLD querier and receives multicast traffic.

To access this page, click **Switching > MLD Snooping > Multicast Router Configuration** in the navigation menu.

To change the multicast router mode for one or more interfaces, select each entry to modify and click **Edit**.

Table 211: Multicast Router Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When configuring the MLD snooping multicast router settings, this field identifies each interface that is being configured.
Multicast Router	Whether the interface is enabled or disabled as a multicast router interface.

Multicast Router VLAN Status

Use the **Multicast Router VLAN Status** page to enable or disable specific VLANs as static multicast router interfaces for a physical port or LAG and to view the multicast router VLAN status for each interface. A multicast router interface faces a multicast router or MLD querier and receives multicast traffic. If a multicast router is attached to the switch, its existence can be learned dynamically. You can also statically configure one or more VLANs on each interface to act as multicast router interfaces.

To access this page, click **Switching > MLD Snooping > Multicast Router VLAN Status** in the navigation menu.

Use the buttons to perform the following tasks:

- To enable one or more VLANs as multicast router interfaces on a port or LAG, click **Add** and configure the settings in the available fields.
- To change the VLANs that are enabled as multicast router interfaces for a port or LAG, select the entry with the settings to change and click **Edit**.
- To disable all VLAN multicast routing interfaces for a port or LAG, select each entry to modify and click **Remove**. You must confirm the action.

Table 212: MLD Snooping Multicast Router VLAN Status Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. Only interfaces that are configured with multicast router VLANs appear in the table. When adding multicast router VLAN information for an interface, use the Interface menu to select the interface on which to enable one or more multicast router VLAN interfaces. When editing multicast router VLAN information, this field shows the interface that is being configured.
VLAN IDs	The ID of each VLAN configured as enabled as a multicast router interface on the associated interface. When changing the multicast routing VLAN interfaces that are associated with an interface, click the VLAN ID to select it (or [Ctrl] + click to select multiple VLAN IDs).

Click **Refresh** to refresh the page with the most current data from the switch.

Configuring MLD Snooping Querier

Use this page to configure the global MLD snooping querier settings on the device. MLD snooping requires that one central switch or router periodically query all end-devices on the network to announce their multicast memberships. This central device is the MLD querier. When Layer 3 IP multicast routing protocols are enabled in a network with IP multicast routing, the IP multicast router acts as the MLD querier. However, if the IP-multicast traffic in a VLAN needs to be Layer 2 switched only, an IP-multicast router is not required. The MLD snooping querier can perform the MLD snooping functions on the VLAN.



Note

This feature is available for 220 switches only.

MLD Snooping Querier Configuration

To access the **MLD Snooping Querier Configuration** page, click **Switching > MLD Snooping Querier > Configuration** in the navigation menu.

Table 213: MLD Snooping Querier Configuration Fields

Field	Description
Admin Mode	The administrative mode for the MLD snooping querier on the device. When enabled, the MLD snooping querier sends out periodic MLD queries that trigger MLD report messages from the switches that want to receive IP multicast traffic. The MLD snooping feature listens to these MLD reports to establish appropriate forwarding.
IPv6 Address	The snooping querier unicast link-local IPv6 address to be used as the source address in periodic MLD queries. This address is used when no IPv6 address is configured on the VLAN on which the query is being sent.
MLD Version	The MLD protocol version used in periodic MLD queries.
Query Interval (Seconds)	The amount of time the MLD snooping querier should wait between sending periodic MLD queries.
Querier Expiry Interval (Seconds)	The amount of time the device remains in non-querier mode after it has discovered that there is a multicast querier in the network.

- If you make any changes to this page, click **Submit** to apply the changes to the system.
- Click **Refresh** to refresh the page with the most current data from the switch.

VLAN Configuration

Use the **MLD Snooping Querier VLAN Configuration** page to enable the MLD snooping querier feature on one or more VLANs and to configure per-VLAN MLD snooping querier settings. Only VLANs that have the MLD snooping querier feature enabled appear in the table.

To access this page, click **Switching > MLD Snooping Querier > VLAN Configuration** in the navigation menu.

Use the buttons to perform the following tasks:

- To enable the MLD snooping querier feature on a VLAN, click **Add** and specify the desired settings.
- To change the MLD snooping querier settings for a VLAN, select the entry to modify and click **Edit**.
- To disable the MLD snooping querier feature on one or more VLANs, select each entry to change and click **Remove**. You must confirm the action before the entry is deleted. Clicking this button does not remove the VLAN from the system.

Table 214: MLD Snooping Querier VLAN Configuration Fields

Field	Description
VLAN ID	The VLAN on which the MLD snooping querier is enabled. When enabling the MLD snooping querier on a VLAN, use this menu to select the desired VLAN. Only VLANs that have been configured on the system and are not already enabled for the MLD snooping querier appear in the menu. When modifying MLD snooping querier settings, this field identifies the VLAN that is being configured.
Querier Election Participation	The participation mode for the MLD snooping querier election process: <ul style="list-style-type: none"> • Enabled: The MLD snooping querier on this VLAN participates in the querier election process when it discovers the presence of another querier in the VLAN. If the snooping querier finds that the other querier source IPv6 address is lower than its own address, it stops sending periodic queries. If the snooping querier wins the election (because it has the lowest IPv6 address), then it continues sending periodic queries. • Disabled: When the MLD snooping querier on this VLAN sees other queriers of the same version in the VLAN, the snooping querier moves to the non-querier state and stops sending periodic queries.
Querier VLAN IPv6 Address	The MLD snooping querier unicast link-local IPv6 address the VLAN uses as the source address in periodic MLD queries sent on the VLAN. If this value is not configured, the VLAN uses the global MLD snooping querier IPv6 address.

Click **Refresh** to refresh the page with the most current data from the switch.

VLAN Status

Use the **MLD Snooping Querier VLAN Status** page to view information about the MLD snooping querier status for all VLANs that have the snooping querier enabled.

To access this page, click **Switching > MLD Snooping Querier > VLAN Status** in the navigation menu.

Table 215: MLD Snooping Querier VLAN Configuration Fields

Field	Description
VLAN ID	The VLAN associated with the rest of the data in the row. The table includes only VLANs that have the snooping querier enabled.
State	The operational state of the MLD Snooping Querier on a VLAN, which is one of the following: <ul style="list-style-type: none"> • Querier: The snooping switch is the querier in the VLAN. The snooping switch will send out periodic queries with a time interval equal to the configured querier query interval. If the snooping switch sees a better querier (numerically lower) in the VLAN, it moves to non-querier mode. • Non-Querier: The snooping switch is in non-querier mode in the VLAN. If the querier expiry interval timer expires, the snooping switch moves into querier mode. • Disabled: The snooping querier is not operational on the VLAN. The snooping querier moves to the disabled mode when MLD snooping is not operational on the VLAN, when the querier address is not configured, or the network management address is not configured.

Table 215: MLD Snooping Querier VLAN Configuration Fields (continued)

Field	Description
Version	The operational MLD protocol version of the querier.
Last IPv6 Address	The IPv6 address of the last querier from which a query was snooped on the VLAN.
Last Version	The MLD protocol version of the last querier from which a query was snooped on the VLAN.
Max Response Time (Seconds)	The maximum response time to be used in the queries that are sent by the snooping querier.

Click **Refresh** to refresh the page with the most current data from the switch.

Creating Port Channels

A port channel, also known as a LAG, allows you to combine multiple full-duplex Ethernet links into a single logical link. Network devices treat the aggregation as if it were a single link, which increases fault tolerance and provides load sharing. You assign the port channel LAG VLAN membership after you create a port channel. The port channel by default becomes a member of the management VLAN.

A port channel (LAG) interface can be either static or dynamic, but not both. All members of a port channel must participate in the same protocols. A static port channel interface does not require a partner system to be able to aggregate its member ports.



Note

If you configure the maximum number of dynamic port channels (LAGs) that your platform supports, additional port channels that you configure are automatically static.

Static LAGs are supported. When a port is added to a LAG as a static member, it neither transmits nor receives LACPDUs.

Port Channel Summary

Use the **Port Channel Summary** page to group one or more full duplex Ethernet links to be aggregated together to form a port channel, which is also known as a LAG. The switch can treat the port channel as if it were a single link.

To access this page, click **Switching > Port Channel > Summary** in the navigation menu.

Table 216: Port Channel Summary Fields

Field	Description
Name	Identifies the user-configured text name of the port channel.
Type	<p>The type of port channel:</p> <ul style="list-style-type: none"> • Dynamic: Uses Link Aggregation Control Protocol (LACP) Protocol Data Units (PDUs) to exchange information with the link partners to help maintain the link state. To utilize Dynamic link aggregation on this port channel, the link partner must also support LACP. • Static: Does not require a partner system to be able to aggregate its member ports. When a port is added to a port channel as a static member, it neither transmits nor receives LACP PDUs. <p>When configuring a port channel, use the Static Mode field to set the port channel type. If the Static Mode is disabled, the port channel type is Dynamic.</p>
Admin Mode	Select enable or disable from the drop-down menu. When the Port Channel is disabled no traffic will flow and LACPDUs will be dropped, but the links that form the Port Channel will not be released. The factory default is enable.
STP Mode	Shows whether the STP Administrative Mode is enabled or disabled on the port channel
Link State	Whether the link is Up or Down.
Link Trap	Shows whether to send traps when link status changes. If the status is Enabled, traps are sent.
Members	Lists the ports that are members of the Port Channel, in Slot/Port notation (Unit/Slot/Port for stackable systems). There can be a maximum of 8 ports assigned to a Port Channel.
Active Ports	Lists the ports that are actively participating members of this Port Channel, in Slot/Port notation (Unit/Slot/Port for stackable systems).
Load Balance	<p>The algorithm used to distribute traffic load among the physical ports of the port channel while preserving the per-flow packet order. The packet attributes the load-balancing algorithm can use to determine the outgoing physical port include the following:</p> <ul style="list-style-type: none"> • Source MAC, VLAN, Ethertype, Incoming Port • Destination MAC, VLAN, Ethertype, Incoming Port • Source/Destination MAC, VLAN, Ethertype, Incoming Port • Source IP and Source TCP/UDP Port Fields • Destination IP and Destination TCP/UDP Port Fields • Source/Destination IP and TCP/UDP Port Fields • Enhanced Hashing Mode

Port Channel Configuration

Use the **Port Channel Configuration** page to group one or more full duplex Ethernet links to be aggregated together to form a port channel, which is also known as a **LAG**. The switch treats the port channel as if it were a single link.

To access this page, click **Switching > Port Channel > Summary** in the navigation menu. Select a port and click **Edit**.

Table 217: Port Channel Configuration Fields

Field	Description
Port Channel Interface	Select the port channel to configure. The port channel follows a Slot/Port (or Unit/Slot/Port for stacking platforms) interface naming convention, where the slot is 3.
Port Channel Name	Enter the name you want assigned to the Port Channel. You may enter any string of up to 15 alphanumeric characters. You must specify a valid name in order to create the Port Channel.
Link Trap	Specify whether you want to have a trap sent when link status changes. The factory default is enable, which will cause the trap to be sent.
Administrative Mode	Select enable or disable from the drop-down menu. When the Port Channel is disabled no traffic will flow and LACPDU's will be dropped, but the links that form the Port Channel will not be released. The factory default is enable.
Link Status	Whether the link is Up or Down.
STP Mode	Select the STP Administrative Mode associated with the Port Channel: <ul style="list-style-type: none"> • Disable: Spanning tree is disabled for this Port Channel. • Enable: Spanning tree is enabled for this Port Channel.
Static Mode	Select enable or disable from the drop-down menu. The factory default is Disable. <ul style="list-style-type: none"> • Enable: The port channel is statically maintained, which means it does not transmit or process received LAGPDUs. The member ports do not transmit LAGPDUs and all the LAGPDUs it may receive are dropped. A static port channel interface does not require a partner system to be able to aggregate its member ports. • Disable: The port channel is dynamically maintained. The interface transmits and processes LAGPDUs and requires a partner system.
Local Preference Mode	This field is available only on systems that support stacking. When this option is enabled, the LAG-destined unicast traffic egresses only out of members of the LAG interface on the local unit. This feature makes sure that the LAG-destined unicast traffic does not cross the external stack link when the LAG has members on the local unit.
Load Balance	Select the hashing algorithm used to distribute the traffic load among available physical ports in the LAG. The range of possible values may vary with the type of switch. The possible values are: <ul style="list-style-type: none"> • Source MAC, VLAN, EtherType, and source port • Destination MAC, VLAN, EtherType and source port • Source/Destination MAC, VLAN, EtherType, and source port • Source IP and Source TCP/UDP Port • Destination IP and Destination TCP/UDP Port • Source/Destination IP and source/destination TCP/UDP Port • Enhanced hashing mode
Port Channel Members	After you create one or more port channel, this field lists the members of the Port Channel. If there are no port channels on the system, this field is not present.

Table 217: Port Channel Configuration Fields (continued)

Field	Description
Unit/Slot/Port	This column lists the physical ports available on the system.
Participation	Select each port's membership status for the Port Channel you are configuring. There can be a maximum of 8 ports assigned to a Port Channel. <ul style="list-style-type: none"> • Include: The port participates in the port channel. • Exclude: The port does not participate in the port channel, which is the default.
Membership Conflicts	Shows ports that are already members of other Port Channels. A port may only be a member of one Port Channel at a time. If the entry is blank, the port is not currently a member of any Port Channel

- If you make any changes to this page, click **Submit** to apply the changes to the system.
- To remove a port channel, select it from the **Port Channel Name** drop-down menu and click delete. All ports that were members of this Port Channel are removed from the Port Channel and included in the default VLAN. This field will not appear when a new Port Channel is being created.

Port Channel Statistics

The **Port Channel Statistics** page displays the flap count for each port channel and their member ports. A flap occurs when a port channel interface or port channel member port goes down.

To access this page, click **Switching > Port Channel > Statistics** in the navigation menu.

Table 218: Port Channel Statistics Fields

Field	Description
Interface	The port channel or member port (physical port) associated with the rest of the data in the row.
Channel Name	The port channel name associated with the port channel. For a physical port, this field identifies the name of the port channel of which the port is a member.
Type	The interface type, which is either Port Channel (logical link-aggregation group) or Member Port (physical port).
Flap Count	The number of times the interface has gone down. The counter for a member port is incremented when the physical port is either manually shut down by the administrator or when its link state is down. When a port channel is administratively shut down, the flap counter for the port channel is incremented, but the flap counters for its member ports are not affected. When all active member ports for a port channel are inactive (either administratively down or link down), then the port channel flap counter is incremented.
Clear Counters (Button)	Click this button to reset the flap counters for all port channels and member ports to 0.

Click **Refresh** to display the latest information from the router.

Viewing Multicast Forwarding Database Information

The Layer 2 Multicast Forwarding Database (MFDB) is used by the switch to make forwarding decisions for packets that arrive with a multicast destination MAC address. By limiting multicasts to only certain ports in the switch, traffic is prevented from going to parts of the network where that traffic is unnecessary.

When a packet enters the switch, the destination MAC address is combined with the VLAN ID and a search is performed in the Layer 2 Multicast Forwarding Database. If no match is found, the packet is either flooded to all ports in the VLAN or discarded, depending on the switch configuration. If a match is found, the packet is forwarded only to the ports that are members of that multicast group.

MFDB Table

Use the **MFDB Table** page to view the port membership information for all active multicast address entries. The key for an entry consists of a VLAN ID and MAC address pair. Entries may contain data for more than one protocol.

To access this page, click **Switching > Multicast Forwarding Database > Summary** in the navigation menu.

Table 219: MFDB Summary Fields

Field	Description
MAC Address	The VLAN ID (the first two groups of hexadecimal digits) and multicast MAC address (the last six groups of hexadecimal digits) that has been added to the MFDB.
Component	The feature on the device that was responsible for adding the entry to the multicast forwarding database, which is one of the following: <ul style="list-style-type: none"> • IGMP Snooping: A Layer 2 feature that allows the device to dynamically add or remove ports from IPv4 multicast groups by listening to <i>IGMP</i> join and leave requests. • MLD Snooping: A Layer 2 feature that allows the device to dynamically add or remove ports from IPv6 multicast groups by listening to MLD join and leave requests. • GMRP: Generic Address Resolution Protocol (GARP) Multicast Registration Protocol, which helps help control the flooding of multicast traffic by keeping track of group membership information. • Static Filtering: A static MAC filter that was manually added to the address table by an administrator.
Type	The type of entry, which is one of the following: <ul style="list-style-type: none"> • Static: The entry has been manually added to the MFDB by an administrator. • Dynamic: The entry has been added to the MFDB as a result of a learning process or protocol.
Description	A text description of this multicast table entry.
Interface(s)	The list of interfaces that will forward or filter traffic sent to the multicast MAC address.
Forwarding Interface(s)	The list of forwarding interfaces. This list does not include any interfaces that are listed as static filtering interfaces.

- To search for a MAC address if the list is too long to scan, enter the MAC address in hex format and click **Search**.
- Click **Refresh** to update the information on the screen with the most current data.

GMRP Table

Use the **Multicast Forwarding Database GMRP Table** page to display the entries in the multicast forwarding database (MFDB) that were added by using the GARP Multicast Registration Protocol (GMRP).

To access this page, click **Switching > Multicast Forwarding Database > GMRP** in the navigation menu.

Table 220: GMRP Fields

Field	Description
VLAN ID	The VLAN ID associated with the entry in the MFDB.
MAC Address	The multicast MAC address associated with the entry in the MFDB.
Type	The type of entry, which is one of the following: <ul style="list-style-type: none"> • Static: The entry has been manually added to the MFDB by an administrator. • Dynamic: The entry has been added to the MFDB as a result of a learning process or protocol. Entries that appear on this page have been added by using GARP.
Description	A text description of this multicast table entry.
Interface(s)	The list of interfaces that will forward or filter traffic sent to the multicast MAC address.

Click **Refresh** to update the information on the screen with the most current data.

IGMP Snooping Table

The **Multicast Forwarding Database IGMP Snooping Table** page displays the entries in the multicast forwarding database (MFDB) that were added because they were discovered by the *IGMP* snooping feature. IGMP snooping allows the device to dynamically add or remove ports from IPv4 multicast groups by listening to IGMP join and leave requests.

To access this page, click **Switching > Multicast Forwarding Database > IGMP Snooping** in the navigation menu.

Table 221: IGMP Snooping Fields

Field	Description
VLAN ID	The VLAN ID associated with the entry in the MFDB.
MAC Address	The multicast MAC address associated with the entry in the MFDB.

Table 221: IGMP Snooping Fields (continued)

Field	Description
Type	The type of entry, which is one of the following: <ul style="list-style-type: none"> • Static: The entry has been manually added to the MFDB by an administrator. • Dynamic: The entry has been added to the MFDB as a result of a learning process or protocol. Entries that appear on this page have been learned by examining IGMP messages.
Description	A text description of this multicast table entry.
Interface(s)	The list of interfaces that will forward or filter traffic sent to the multicast MAC address.

- Click **Refresh** to update the information on the screen with the most current data.

MFDB Statistics

Use the **Multicast Forwarding Database Statistics** page to view statistical information about the MFDB table.

To access this page, click **Switching > Multicast Forwarding Database > Statistics** in the navigation menu.

Table 222: Multicast Forwarding Database Statistics Fields

Field	Description
MFDB Max Table Entries	The maximum number of entries that the multicast forwarding database can hold.
MFDB Most Entries Since Last Reset	The largest number of entries that have been present in the multicast forwarding database since the switch was last rebooted. This value is also known as the MFDB high-water mark.
MFDB Current Entries	The current number of entries in the multicast forwarding database.

Click **Refresh** to update the information on the screen with the most current data.

Multicast VLAN Registration

Multicast VLAN Registration (MVR) allows the switch to listen to the *IGMP* frames. Both protocols operate independently from each other and can be enabled on the switch interfaces. In such case, MVR listens to the Join and Report messages only for the statically configured groups. All other groups are managed by IGMP snooping. MVR uses the multicast VLAN, a dedicated VLAN used to transfer multicast traffic over the network avoiding duplication of multicast streams for clients in different VLANs.

MVR Global Configuration

Use the **MVR Global Configuration** page to view and configure the global settings for MVR.

To access this page, click **Switching > MVR > Global** in the navigation menu.

Table 223: MVR Global Configuration Fields

Field	Description
Admin Mode	The administrative mode of MVR on the device.
MVR Mode	<p>The MVR learning mode, which can be one of the following:</p> <ul style="list-style-type: none"> • Compatible: MVR does not learn source ports membership; instead, all source ports are members of all groups by default. MVR does not forward <i>IGMP</i> Joins and Leaves from the hosts to the router. • Dynamic: MVR learns source ports membership from IGMP queries. MVR forwards the IGMP Joins and Leaves from the hosts to the router. <p>The multicast traffic is forwarded only to the receiver ports that joined the group, either by IGMP Joins or MVR static configuration.</p>
Multicast VLAN	A dedicated VLAN used to transfer multicast traffic over the network, avoiding duplication of multicast streams for clients in different VLANs.
Maximum Multicast Groups	The maximum number of membership groups that can be statically configured in the MVR database.
Current Multicast Groups	The current number of membership groups that are statically configured in the MVR database.
Query Response Time	The maximum time to wait for an IGMP membership report on a receiver port before removing the port from the multicast group. The query time is specified in tenths of a second.

- If you make any configuration changes, click **Submit** to apply the new settings to the switch.
- Click **Refresh** to update the information on the screen with the most current data.

MVR Group Status

Use the **MVR Group Status** page to view or configure MVR groups. MVR maintains two types of group entries in its database, Static and Dynamic. Static entries are configured by the administrator and dynamic entries are learned by MVR on the source ports.

To access this page, click **Switching > MVR > Group** in the navigation menu.

Use the buttons to perform the following tasks:

- To add a group, click **Add** and specify a group address in the available field.
- To edit a configured group, select the entry to modify and click **Edit**. Then, configure which interfaces should be members of that group.
- To remove one or more configured groups, select each entry to delete and click Remove. You must confirm the action before the entry is deleted.

Table 224: MVR Group Status Fields

Field	Description
Group	The multicast group address.
Status	The status of the group, which can be one of the following: <ul style="list-style-type: none"> • Active: Group has one or more MVR ports participating. • Inactive: Group has no MVR ports participating.
Members	The list of interfaces which participate in the MVR group. In the compatible mode, all source ports are members of all groups by default.
Contiguous Group Count	This field is available in the Add Group dialog. Specify the desired number of groups to be created starting with the entered group address. The default contiguous group count is 1.
Available Interfaces	This field is available in the Edit Group Configuration dialog. The interfaces that can be added to the group. To move an interface between the Available Interfaces and Group Interfaces fields, click the interface (or [Ctrl] + click to select multiple interfaces), and then click the appropriate arrow to move the selected interfaces to the desired field.
Group Interfaces	This field is available in the Edit Group Configuration dialog. The interfaces that are members of the MVR group.

Click **Refresh** to update the information on the screen with the most current data.

MVR Interface Status

Use the **MVR Interface Status** page to configure MVR settings on specific interfaces. To configure the settings for one or more interfaces, select each entry to modify and click **Edit**. The same MVR settings are applied to all selected interfaces.

To access this page, click **Switching > MVR > Interface** in the navigation menu.

Table 225: MVR Interface Status Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When configuring MVR settings, this field identifies the interface(s) that are being configured.
MVR Interface Mode	The administrative mode of MVR on the interface. MVR must be enabled globally and on an interface in order to listen to the Join and Report messages for the configured groups.
Type	The type of interface, which can be one of the following: <ul style="list-style-type: none"> • Source: The port where multicast traffic is flowing to. It must be a member of the multicast VLAN. • Receiver: The port where listening host is connected to the switch. It must not be a member of the multicast VLAN. • None: The port is not an MVR port.

Table 225: MVR Interface Status Fields (continued)

Field	Description
Status	The active state of the interface, which can be one of the following: <ul style="list-style-type: none"> • Active: The port has link up and is in the forwarding state. • Inactive: The port may not have link up, not be in the forwarding state, or both.
Immediate Leave	The MVR immediate leave mode on the interface. It can only be configured on the receiver ports. MVR handles <i>IGMP</i> Leaves in Normal or Immediate leave mode. When a Leave message is received, in the normal mode a general IGMP query is sent from the switch to the receiver port, whereas in the immediate leave mode the switch is immediately reconfigured not to forward specific multicast stream to the receiver port. The immediate leave mode is used for ports where only one client may be connected.

Click **Refresh** to update the information on the screen with the most current data.

MVR Statistics

Use the **MVR Statistics** page to view statistical information about *IGMP* packets intercepted by MVR.

To access this page, click **Switching > MVR > Statistics** in the navigation menu.

Table 226: MVR Statistics Fields

Field	Description
IGMP Queries	The total number of IGMP Queries successfully transmitted or received by the processor.
IGMPv1 Reports	The total number of IGMPv1 Reports successfully transmitted or received by the processor.
IGMPv2 Reports	The total number of IGMPv2 Reports successfully transmitted or received by the processor.
IGMP Leaves	The total number of IGMP Leaves successfully transmitted or received by the processor.
Packet Failures	The total number of packets which failed to get transmitted or received by the processor.

Click **Refresh** to update the information on the screen with the most current data.

Configuring Protected Ports

Use the **Protected Ports Configuration** page to configure and view protected ports groups. A port that is a member of a protected ports group is a protected port. A port that is not a member of any protected ports group is an unprotected port. Each port can be a member of only one protected ports group. Ports in the same protected ports group cannot forward traffic to other protected ports within the group, even if they are members of the same VLAN. However, a port in a protected ports group can forward traffic to ports that are in a different protected ports group. A protected port can also forward traffic to unprotected ports. Unprotected ports can forward traffic to both protected and unprotected ports.

To access this page, click **Switching > Protected Ports > Configuration** in the navigation menu.

Use the buttons to perform the following tasks:

- To create a protected ports group and add ports to the group, click **Add** and configure the settings in the available fields.
- To change the name or the port members for an existing group, select the group to update and click **Edit**.
- To remove one or more protected ports groups, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.

Table 227: Protected Ports Configuration Fields

Field	Description
Group Name	This is the configured name of the protected ports group.
Protected Ports	The ports that are members of the protected ports group. When adding a port to a protected ports group, the Available Interfaces field lists the ports that are not already members of a protected ports group. To move an interface between the Available Interfaces and Selected Interfaces fields, click the port (or [Ctrl] + click to select multiple ports), and then click the appropriate arrow to move the port(s) to the desired field.

Click **Refresh** to update the information on the screen with the most current data.

Configuring Spanning Tree Protocol

The *STP (Spanning Tree Protocol)* provides a tree topology for any arrangement of bridges. STP also provides one path between end stations on a network, eliminating loops. Spanning tree versions supported include Common STP, Multiple STP, and Rapid STP.

Classic STP provides a single path between end stations, avoiding and eliminating loops. For information on configuring Common STP, see [CST Port Configuration](#) on page 227.

MSTP (Multiple Spanning Tree Protocol) supports multiple instances of Spanning Tree to efficiently channel VLAN traffic over different interfaces. Each instance of the Spanning Tree behaves in the manner specified in IEEE 802.1w, Rapid Spanning Tree (RSTP), with slight modifications in the working but not the end effect (chief among the effects, is the rapid transitioning of the port to 'Forwarding'). The difference between the RSTP and the traditional STP (IEEE 802.1D) is the ability to configure and recognize full duplex connectivity and ports which are connected to end stations, resulting in rapid transitioning of the port to 'Forwarding' state and the suppression of Topology Change Notification. These features are represented by the parameters 'pointtopoint' and 'edgeport'. MSTP is compatible to both RSTP and STP. It behaves appropriately to STP and RSTP bridges. A MSTP bridge can be configured to behave entirely as a RSTP bridge or a STP bridge.



Note

For two bridges to be in the same region, the force version should be 802.1S and their configuration name, digest key, and revision level should match. For more information about regions and their effect on network topology, refer to the IEEE 802.1Q standard.

Switch Configuration/Status

The **Spanning Tree Switch Configuration/Status** page contains fields for enabling *STP* on the switch.

To access this page, click **Switching > Spanning Tree > Switch** in the navigation menu.

Table 228: Spanning Tree Switch Configuration Fields

Field	Description
Spanning Tree Admin Mode	The administrative mode of STP on the device. When enabled, the device participates in the root bridge election process and exchanges Bridge Protocol Data Units (BPDUs) with other switches in the spanning tree to determine the root path costs and maintain topology information.
Force Protocol Version	The STP version the device uses, which is one of the following: <ul style="list-style-type: none"> • IEEE 802.1d: Classic STP provides a single path between end stations, avoiding and eliminating loops. • IEEE 802.1w: Rapid Spanning Tree Protocol (RSTP) behaves like classic STP but also can configure and recognize full-duplex connectivity and ports that are connected to end stations, resulting in rapid transitioning of the port to the Forwarding state and the suppression of Topology Change Notifications. • IEEE 802.1s: <i>MSTP</i> includes all the advantages of RSTP and also supports multiple spanning tree instances to efficiently channel VLAN traffic over different interfaces. MSTP is compatible with both RSTP and STP.
Configuration Name	The name of the MSTP region. Each switch that participates in the same MSTP region must share the same Configuration Name, Configuration Revision Level, and MST-to-VLAN mappings.
Configuration Revision Level	The revision number of the MSTP region. This number must be the same on all switches that participate in the MSTP region.
Configuration Digest Key	The 16 byte signature of type HMAC-MD5 (<i>Message-Digest algorithm 5</i>) created from the MST Configuration Table (a VLAN ID-to-MST ID mapping).
Configuration Format Selector	The version of the configuration format being used in the exchange of BPDUs.

- If you make any configuration changes, click **Submit** to apply the new settings to the switch.
- Click **Refresh** to update the information on the screen with the most current data.

CST Configuration

Use the **CST Configuration** page to configure the Common Spanning Tree (CST) settings. The settings and information on this page define the device within the spanning tree topology that connects all *STP*/RSTP bridges and *MSTP* regions.

To access this page, click **Switching > Spanning Tree > CST** in the navigation menu.

Table 229: Spanning Tree CST Fields

Field	Description
Bridge Priority	The value that helps determine which bridge in the spanning tree is elected as the root bridge during STP convergence. A lower value increases the probability that the bridge becomes the root bridge.
Bridge Max Age	The amount of time a bridge waits before implementing a topological change.
Bridge Hello Time	The amount of time the root bridge waits between sending hello BPDUs.

Table 229: Spanning Tree CST Fields (continued)

Field	Description
Bridge Forward Delay	The amount of time a bridge remains in a listening and learning state before forwarding packets.
Spanning Tree Maximum Hops	The maximum number of hops a Bridge Protocol Data Unit (BPDU) is allowed to traverse within the spanning tree region before it is discarded.
BPDU Guard	When enabled, BPDU Guard can disable edge ports that receive BPDU packets. This prevents a new device from entering the existing STP topology. Thus devices that were originally not a part of STP are not allowed to influence the STP topology.
BPDU Filter	When enabled, this feature filters the BPDU traffic on the edge ports. When spanning tree is disabled on a port, BPDU filtering allows BPDU packets received on that port to be dropped.
Spanning Tree Tx Hold Count	The maximum number of BPDUs that a bridge is allowed to send within a hello time window.
Bridge Identifier	A unique value that is automatically generated based on the bridge priority value and the base MAC address of the bridge. When electing the root bridge for the spanning tree, if the bridge priorities for multiple bridges are equal, the bridge with the lowest MAC address is elected as the root bridge.
Time Since Topology Change	The amount of time that has passed since the topology of the spanning tree has changed.
Topology Change Count	The number of times the topology of the spanning tree has changed.
Topology Change	Whether a topology change is in progress on any port assigned to the CST. If a change is in progress the value is True; otherwise, it is False.
Designated Root	The bridge identifier of the root bridge for the CST. The identifier is made up of the bridge priority and the base MAC address.
Root Path Cost	The path cost to the designated root for the CST. Traffic from a connected device to the root bridge takes the least-cost path to the bridge. If the value is 0, the cost is automatically calculated based on port speed.
Root Port	The port on the bridge with the least-cost path to the designated root for the CST.
Max Age	The amount of time a bridge waits before implementing a topological change.
Forward Delay	The forward delay value for the root port bridge.
Hold Time	The minimum amount of time between transmissions of Configuration BPDUs.
CST Regional Root	The bridge identifier of the CST regional root. The identifier is made up of the priority value and the base MAC address of the regional root bridge.
CST Path Cost	The path cost to the CST tree regional root.

- If you make any configuration changes, click **Submit** to apply the new settings to the switch.
- Click **Force** to force the port to send out 802.1w or 802.1D BPDUs.
- Click **Refresh** to update the screen with most recent data.

CST Port Configuration

Use the **Spanning Tree CST Port Configuration/Status** page to view and configure the Common Spanning Tree (CST) settings for each interface on the device. To configure CST settings for an interface and to view additional information about the interface's role in the CST topology, select the interface to view or configure and click **Edit**.

To access this page, click **Switching > Spanning Tree > CST Port** in the navigation menu.

Table 230: Spanning Tree CST Port Fields

Field	Description
Interface	The port or <i>LAG</i> associated with the rest of the data in the row. When configuring CST settings for an interface, this field identifies the interface being configured.
Port Role	The role of the port within the CST, which is one of the following: <ul style="list-style-type: none"> • Root: A port on the non-root bridge that has the least-cost path to the root bridge. • Designated: A port that has the least-cost path to the root bridge on its segment. • Alternate: A blocked port that has an alternate path to the root bridge. • Backup: A blocked port that has a redundant path to the same network segment as another port on the bridge. • Master: The port on a bridge within an MST instance that links the MST instance to other <i>STP</i> regions. • Disabled: The port is administratively disabled and is not part of the spanning tree.
Port Forwarding State	<ul style="list-style-type: none"> • Blocking: The port discards user traffic and receives, but does not send, BPDUs. During the election process, all ports are in the blocking state. The port is blocked to prevent network loops. • Listening: The port sends and receives BPDUs and evaluates information to provide a loop-free topology. This state occurs during network convergence and is the first state in transitioning to the forwarding state. • Learning: The port learns the MAC addresses of frames it receives and begins to populate the MAC address table. This state occurs during network convergence and is the second state in transitioning to the forwarding state. • Forwarding: The port sends and receives user traffic. • Disabled: The port is administratively disabled and is not part of the spanning tree.
Port Priority	The priority for the port within the CST. This value is used in determining which port on a switch becomes the root port when two ports have the same least-cost path to the root. The port with the lower priority value becomes the root port. If the priority values are the same, the port with the lower interface index becomes the root port.
Port Path Cost	The path cost from the port to the root bridge.
Description	A user-configured description of the port. After you select an interface and click Edit , a window opens and allows you to edit the CST port settings and view additional CST information for the interface. The following information describes the additional fields available in the Edit CST Port Entry window.

Table 230: Spanning Tree CST Port Fields (continued)

Field	Description
Admin Edge Port	Select this option administratively configure the interface as an edge port. An edge port is an interface that is directly connected to a host and is not at risk of causing a loop.
Auto-calculate Port Path Cost	Shows whether the path cost from the port to the root bridge is automatically determined by the speed of the interface (Enabled) or configured manually (Disabled).
Hello Timer	The amount of time the port waits between sending hello BPDUs.
External Port Path Cost	The cost of the path from the port to the CIST root. This value becomes important when the network includes multiple regions.
Auto-calculate External Port Path Cost	Shows whether the path cost from the port to the CIST root is automatically determined by the speed of the interface (Enabled) or configured manually (Disabled).
BPDU Flood	This option determines the behavior of the interface if STP is disabled on the port and the port receives a BPDU. If BPDU flooding is enabled, the port will flood the received BPDU to all the ports on the switch that are similarly disabled for spanning tree.
BPDU Guard Effect	Shows the status of BPDU Guard Effect on the interface. When enabled, BPDU Guard Effect can disable edge ports that receive BPDU packets. This prevents a new device from entering the existing STP topology. Thus devices that were originally not a part of STP are not allowed to influence the STP topology.
Port ID	A unique value that is automatically generated based on the port priority value and the interface index.
Port Up Time Since Counters Last Cleared	The amount of time that the port has been up since the counters were cleared.
Port Mode	The administrative mode of spanning tree on the port.
Designated Root	The bridge ID of the root bridge for the CST.
Designated Cost	The path cost offered to the LAN by the designated port.
Designated Bridge	The bridge ID of the bridge with the designated port.
Designated Port	The port ID of the designated port.
Topology Change Acknowledge	Whether the next BPDU to be transmitted for this port will have the topology change acknowledgement flag set.
Auto Edge	When enabled, Auto Edge allows the interface to become an edge port if it does not receive any BPDUs within a given amount of time.
Edge Port	Whether the interface is configured as an edge port (Enabled).
Point-to-point MAC	Whether the link type for the interface is a point-to-point link.
Root Guard	When enabled, Root Guard allows the interface to discard any superior information it receives to protect the root of the device from changing. The port gets put into discarding state and does not forward any frames.
Loop Guard	When enabled, Loop Guard prevents an interface from erroneously transitioning from blocking state to forwarding when the interface stops receiving BPDUs. The port is marked as being in loop-inconsistent state. In this state, the interface does not forward frames.

Table 230: Spanning Tree CST Port Fields (continued)

Field	Description
TCN Guard	When enabled, TCN Guard restricts the interface from propagating any topology change information received through that interface.
CST Regional Root	The bridge ID of the bridge that has been elected as the root bridge of the CST region.
CST Path Cost	The path cost from the interface to the CST regional root.
Loop Inconsistent State	Identifies whether the interface is currently in a loop inconsistent state. An interface transitions to a loop inconsistent state if loop guard is enabled and the port stops receiving BPDUs. In this state, the interface does not transmit frames.
Transitions Into LoopInconsistent State	The number of times this interface has transitioned into loop inconsistent state.
Transitions Out Of LoopInconsistent State	The number of times this interface has transitioned out of loop inconsistent state.

- If you make any configuration changes, click **Submit** to apply the new settings to the switch.
- Click **Cancel** to cancel the change.
- Click **Refresh** to update the screen with most recent data.

MST Configuration

Use the **MST Configuration** page to view and configure each *MSTI (Multiple Spanning Tree Instances)* on the device. *MSTP* allows the creation of MSTIs based upon a VLAN or groups of VLANs. Configuring MSTIs creates an active topology with a better distribution of network traffic and an increase in available bandwidth when compared to classic STP.

- Use the buttons to perform the following tasks:
- To configure a new MSTI, click **Add** and specify the desired settings.
- To change the Priority or the VLAN associations for an existing MSTI, select the entry to modify and click **Edit**.
- To remove one or more MSTIs, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.

To access this page, click **Switching > Spanning Tree > MST** in the navigation menu.

Table 231: Spanning Tree MST Summary Fields

Field	Description
MST ID	The number that identifies the MST instance.
Priority	The bridge priority for the spanning-tree instance. This value affects the likelihood that the bridge is selected as the <i>root bridge</i> . A lower value increases the probability that the bridge is selected as the root bridge.
# of Associated VLANs	The number of VLANs that are mapped to the MSTI. This number does not contain any information about the VLAN IDs that are mapped to the instance.

Table 231: Spanning Tree MST Summary Fields (continued)

Field	Description
Bridge Identifier	A unique value that is automatically generated based on the bridge priority value of the MSTI and the base MAC address of the bridge. When electing the root bridge for an MST instance, if the bridge priorities for multiple bridges are equal, the bridge with the lowest MAC address is elected as the root bridge.
Time Since Topology Change	The amount of time that has passed since the topology of the MSTI has changed.
Designated Root	The bridge identifier of the root bridge for the MST instance. The identifier is made up of the bridge priority and the base MAC address.
Root Path Cost	The path cost to the designated root for this MST instance. Traffic from a connected device to the root bridge takes the least-cost path to the bridge. If the value is 0, the cost is automatically calculated based on port speed.
Root Port	The port on the bridge with the least-cost path to the designated root for the MST instance.

- If you make any configuration changes, click **Submit** to apply the new settings to the switch.
- Click **Refresh** to update the screen with most recent data.

MST Port Configuration

Use **MST Port Configuration** page to view and configure the Multiple Spanning Tree (MST) settings for each interface on the device. To configure MST settings for an interface and to view additional information about the interface's role in the MST topology, first select the appropriate MST instance from the MST ID menu. Then, select the interface to view or configure and click **Edit**.

To access this page, click **Switching > Spanning Tree > MST Port** in the navigation menu.



Note

If no MST instances have been configured on the switch, the page displays a **No MSTs Available** message and does not display the fields shown in [Table 232](#).

Table 232: Spanning Tree MST Port Configuration Fields

Field	Description
MST ID	The menu contains the ID of each MST instance that has been created on the device.
Interface	The port or <i>LAG</i> associated with the rest of the data in the row. When configuring MST settings for an interface, this field identifies the interface being configured.

Table 232: Spanning Tree MST Port Configuration Fields (continued)

Field	Description
Port Role	The role of the port within the MST, which is one of the following: <ul style="list-style-type: none"> • Root: A port on the non-root bridge that has the least-cost path to the <i>root bridge</i>. • Designated: A port that has the least-cost path to the root bridge on its segment. • Alternate: A blocked port that has an alternate path to the root bridge. • Backup: A blocked port that has a redundant path to the same network segment as another port on the bridge. • Master: The port on a bridge within an MST instance that links the MST instance to other <i>STP</i> regions. • Disabled: The port is administratively disabled and is not part of the spanning tree.
Port Forwarding State	<ul style="list-style-type: none"> • Blocking: The port discards user traffic and receives, but does not send, <i>BPDUs (Bridge Protocol Data Unit)</i> units. During the election process, all ports are in the blocking state. The port is blocked to prevent network loops. • Listening: The port sends and receives BPDUs and evaluates information to provide a loop-free topology. This state occurs during network convergence and is the first state in transitioning to the forwarding state. • Learning: The port learns the MAC addresses of frames it receives and begins to populate the MAC address table. This state occurs during network convergence and is the second state in transitioning to the forwarding state. • Forwarding: The port sends and receives user traffic. • Disabled: The port is administratively disabled and is not part of the spanning tree.
Port Priority	The priority for the port within the <i>MSTI</i> . This value is used in determining which port on a switch becomes the root port when two ports have the same least-cost path to the root. The port with the lower priority value becomes the root port. If the priority values are the same, the port with the lower interface index becomes the root port.
Port Path Cost	The path cost from the port to the root bridge.
Description	A user-configured description of the port. After you select an interface and click Edit , a window opens and allows you to edit the MST port settings and view additional MST information for the interface. The following information describes the additional fields available in this window.
Auto-calculate Port Path Cost	Shows whether the path cost from the port to the root bridge is automatically determined by the speed of the interface (Enabled) or configured manually (Disabled).
Port ID	A unique value that is automatically generated based on the port priority value and the interface index.
Port Up Time Since Counters Last Cleared	The amount of time that the port has been up since the counters were cleared.
Port Mode	The administrative mode of spanning tree on the port.
Designated Root	The bridge ID of the root bridge for the MST instance.
Designated Cost	The path cost offered to the LAN by the designated port.
Designated Bridge	The bridge ID of the bridge with the designated port.

Table 232: Spanning Tree MST Port Configuration Fields (continued)

Field	Description
Designated Port	The port ID of the designated port.
Loop Inconsistent State	Identifies whether the interface is currently in a loop inconsistent state. An interface transitions to a loop inconsistent state if loop guard is enabled and the port stops receiving BPDUs. In this state, the interface does not transmit frames.
Transitions Into LoopInconsistent State	The number of times this interface has transitioned into loop inconsistent state.
Transitions Out Of LoopInconsistent State	The number of times this interface has transitioned out of loop inconsistent state.

- If you make any configuration changes, click **Submit** to apply the new settings to the switch.
- Click **Refresh** to update the screen with most recent data.

Spanning Tree Statistics

Use the **Spanning Tree Statistics** page to view information about the number and type of bridge protocol data units (BPDUs) transmitted and received on each port.

To access this page, click **Switching > Spanning Tree > Statistics** in the navigation menu.

Table 233: Spanning Tree Statistics Fields

Field	Description
Interface	The port or <u>LAG</u> associated with the rest of the data in the row.
STP BPDUs Rx	The number of classic <u>STP</u> (IEEE 802.1d) BPDUs received by the interface.
STP BPDUs Tx	The number of classic STP BPDUs sent by the interface.
RSTP BPDUs Rx	The number of RSTP (IEEE 802.1w) BPDUs received by the interface.
RSTP BPDUs Tx	The number of RSTP BPDUs sent by the interface.
MSTP BPDUs Rx	The number of <u>MSTP</u> (IEEE 802.1s) BPDUs received by the interface.
MSTP BPDUs Tx	The number of MSTP BPDUs sent by the interface.

- Click **Refresh** to update the screen with most recent data.

PVST Global

Use the **PVST Global** page to view and configure Per VLAN Spanning Tree Protocol (PVSTP)/Per VLAN Rapid Spanning Tree Protocol (PVRSTP) Global settings for the device.

To access this page, click **Switching > Spanning Tree > PVST Global** in the navigation menu.

Table 234: PVSTP/PVRSTP Global Fields

Field	Description
Status	PVSTP/PVRSTP configuration operational mode.
Fast Backbone	Configures Fast Backbone mode. When enabled, the switch detects the indirect link failures and accelerates the spanning tree convergence.
Fast Uplink	Configures Fast Uplink mode.
Max Update Rate (pps)	Configures Fast Uplink's Maximum Update Rate.

- If you make any configuration changes, click **Submit** to apply the new settings to the switch.
- Click **Refresh** to update the screen with most recent data.

PVST VLAN

Use the **PVST VLAN** page to view and configure Per VLAN Spanning Tree Protocol (PVSTP)/Per VLAN Rapid Spanning Tree Protocol (PVRSTP) VLAN settings for the device.

To access this page, click **Switching > Spanning Tree > PVST VLAN** in the navigation menu.

Table 235: PVSTP/PVRSTP VLAN Details Fields

Field	Description
VLAN ID	The unique VLAN identifier (VID).
Root	Configures the switch to become the root bridge or standby root bridge by modifying the bridge priority to a lower value to ensure the bridge is the root (or standby) bridge.
Hello Time (Seconds)	The interval between sending successive BDPUs. Configures the spanning tree hello time interval for the specified VLAN.
Forward Delay (Seconds)	Configures the spanning tree forward delay time for a specified VLAN. This interval is the time spent in listening and learning states before transitioning a port to the forwarding states.
Max Age (Seconds)	The maximum age time before a bridge port saves its configuration information.
Bridge Priority	Configures the bridge priority of a VLAN. The value will be automatically adjusted (rounded) as needed by the system. If the value configured is not among the specified values, then it will be rounded off to the nearest valid value.
To view details of any VLAN, the entry needs to be selected and Details button need to be pressed.	
Root ID	
Priority	The root ID priority for the specified VLAN.
Address	The root ID MAC address for the specified VLAN.
Cost	The root ID cost for the specified VLAN.
Port	The root ID port for the specified VLAN.
Hello Time (Seconds)	The root ID hello time for the specified VLAN.
Max Age (Seconds)	The maximum age for the specified VLAN.

Table 235: PVSTP/PVRSTP VLAN Details Fields (continued)

Field	Description
Forward Delay (Seconds)	The root ID forward delay for the specified VLAN.
Bridge ID	
Priority	The bridge ID priority for the specified VLAN.
Address	The bridge ID MAC address for the specified VLAN.
Hello Time (Seconds)	The bridge ID hello time for the specified VLAN.
Max Age (Seconds)	The bridge ID maximum age for the specified VLAN.
Forward Delay (Seconds)	The bridge ID forward delay for the specified VLAN.
Aging Time (Seconds)	The bridge ID aging time for the specified VLAN.
Interface Details	
Interface	Interface which participates in the specified VLAN.
Role	The role of the interface.
Status	The status of the interface.
Cost	The cost value of the interface.
Prio.Nbr	The priority and neighbor of the interface.

Table 236: PVSTP/PVRSTP VLAN Add/Edit Fields

Field	Description
Root	Configures the switch to become the root bridge or standby root bridge by modifying the bridge priority to a lower value to ensure the bridge is the root (or standby) bridge.
Hello Time (Seconds)	The interval between sending successive BDPUs. Configures the spanning tree hello time interval for the specified VLAN.
Forward Delay (Seconds)	Configures the spanning tree forward delay time for a specified VLAN. This interval is the time spent in listening and learning states before transitioning a port to the forwarding states.
Max Age (Seconds)	The maximum age time before a bridge port saves its configuration information.
Bridge Priority	Configures the bridge priority of a VLAN. The value will be automatically adjusted (rounded) as needed by the system. If the value configured is not among the specified values, then it will be rounded off to the nearest valid value.

- Click **Refresh** to update the screen with most recent data.
- Click **Add** to add a new row to the VLAN configuration
- Select an entry and then click **Edit** to change the PVST configuration on the VLAN.
- Select an entry and then click **Remove** to remove the PVST row from the VLAN configuration.

PVST Interface

Use the **PVST Interface** page to view and configure Per VLAN Spanning Tree Protocol (PVSTP)/Per VLAN Rapid Spanning Tree Protocol (PVRSTP) Interface settings for the device.

To access this page, click **Switching > Spanning Tree > PVST Interface** in the navigation menu.

Table 237: PVSTP/PVRSTP Interface Fields

Field	Description
Interface	List of physical interfaces and LAGs.
Priority	The port priority configuration is used to allow the operator to select the relative importance of the port in the selection process for forwarding. Set this value to a lower number to prefer a port for forwarding of frames. This field is available for configuration only when PVSTP/PVRSTP is enabled.
Per VLAN Configuration	Configuration of each VLAN.
VLAN ID	The unique VLAN identifier (VID).
Priority	The per VLAN priority value configuration of the port is the priority used to allow the operator to select the relative importance of the port in the selection process for forwarding. Set this value to a lower number to prefer a port for forwarding of frames. This priority configuration is used when the port is configured as a point-to-point link type.
Cost	The path cost from the port to the root bridge.

Table 238: PVSTP/PVRSTP Interface Edit Fields

Field	Description
Interface	List of physical interfaces and LAGs.
Priority	The port priority configuration is used to allow the operator to select the relative importance of the port in the selection process for forwarding. Set this value to a lower number to prefer a port for forwarding of frames. This field is available for configuration only when PVSTP/PVRSTP is enabled.
Cost	The path cost from the port to the root bridge.

- Click **Refresh** to update the screen with most recent data.
- Select an entry and then click **Edit** to change the PVST interface configuration.

PVST Statistics

Use the **PVST Statistics** page to view and configure Per VLAN Spanning Tree Protocol (PVSTP)/Per VLAN Rapid Spanning Tree Protocol (PVRSTP) Statistics settings for the device.

To access this page, click **Switching > Spanning Tree > PVST Statistics** in the navigation menu.

Table 239: PVSTP/PVRSTP Statistics Fields

Field	Description
Fast Backbone	
Transition via Fast Backbone	Number of fast backbone transitions.
Inferior BPDUs Received	Number of the received inferior BPDUs.
RLQ Request PDUs Received	Number of the received RLQ request PDUs.
RLQ Response PDUs Received	Number of the received RLQ response PDUs.
RLQ Request PDUs Sent	Number of the sent RLQ request PDUs.
RLQ Response PDUs Sent	Number of the sent RLQ response PDUs.
Fast Uplink	
Fast Uplink Transitions	Number of the fast uplink transitions.
Proxy Multicast Addresses Transmitted	Number of the transmitted proxy multicast addresses.

Click **Refresh** to update the screen with most recent data.

Mapping 802.1p Priority

The IEEE 802.1p feature allows traffic prioritization at the MAC level. The switch can prioritize traffic based on the 802.1p tag attached to the Layer 2 frame. Each port on the switch has multiple queues to give preference to certain packets over others based on the CoS criteria you specify. When a packet is queued for transmission in a port, the rate at which it is serviced depends on how the queue is configured and possibly the amount of traffic present in the other queues of the port. If a delay is necessary, packets get held in the queue until the scheduler authorizes the queue for transmission.

Use the **802.1p Priority Mapping** page in the Class of Service folder to assign 802.1p priority values to various traffic classes on one or more interfaces.

To access this page, click **Switching > Class of Service > 802.1p** in the navigation menu.

Table 240: 802.1p Priority Mapping

Field	Description
Interface	The interface associated with the rest of the data in the row. The Global entry represents the common settings for all interfaces, unless specifically overridden individually.
Priority	The heading row lists each 802.1p priority value (0-7), and the data in the table shows which traffic class is mapped to the priority value. Incoming frames containing the designated 802.1p priority value are mapped to the corresponding traffic class in the device.

Table 240: 802.1p Priority Mapping (continued)

Field	Description
802.1p Priority	The 802.1p priority value to be mapped.
Traffic Class	The internal traffic class to which the corresponding 802.1p priority value is mapped. The default value for each 802.1p priority level is displayed for reference.

Configuring Port Security

Port Security can be enabled on a per-port basis. When a port is locked, only packets with allowable source MAC addresses can be forwarded. All other packets are discarded. A MAC address can be defined as allowable by one of two methods: dynamically or statically. Note that both methods are used concurrently when a port is locked.

Dynamic locking implements a “first arrival” mechanism for Port Security. You specify how many addresses can be learned on the locked port. If the limit has not been reached, a packet with an unknown source MAC address is learned and forwarded normally. Once the limit is reached, no more addresses are learned on the port. Any packets with source MAC addresses that were not already learned are discarded. Note that you can effectively disable dynamic locking by setting the number of allowable dynamic entries to zero.

Static locking allows you to specify a list of MAC addresses that are allowed on a port. The behavior of packets is the same as for dynamic locking: only packets with an allowable source MAC address can be forwarded.

To see the MAC addresses learned on a specific port, see [Configuring and Searching the Forwarding Database](#) on page 121.

Disabled ports can only be activated from the [Configuring Ports](#) page.

Port Security Administration

Use the [Port Security Global Administration](#) page to enable or disable the port security feature on your switch.

To access this page, click **Switching > Port Security > Global** in the navigation menu.

Select **Enable** or **Disable** from the Port Security Mode list and click **Submit**.

Port Security Interface Configuration

Use the [Port Security Interface Status](#) page to configure the port security feature on a selected interface.

To access this page, click **Switching > Port Security > Interface** in the navigation menu.

Use the buttons as follows:

- To configure the settings for one or more interfaces, select each entry to modify and click **Edit**.

- o apply the same settings to all interfaces, click **Edit All**.
- Click **Submit** to apply the new settings to the system.

Table 241: Port Security Interface Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When configuring the port security settings for one or more interfaces, this field lists the interfaces that are being configured.
Port Security Mode	The administrative mode of the port security feature on the interface. The port security mode must be enabled both globally and on an interface to enforce the configured limits for the number of static and dynamic MAC addresses allowed on that interface.
Max Dynamic Addresses Allowed	The number of source MAC addresses that can be dynamically learned on an interface. If an interface reaches the configured limit, any other addresses beyond that limit are not learned, and the frames are discarded. Frames with a source MAC address that has already been learned will be forwarded. A dynamically-learned MAC address is removed from the MAC address table if the entry ages out, the link goes down, or the system reboots. Note that the behavior of a dynamically-learned address changes if the sticky mode for the interface is enabled or the address is converted to a static MAC address.
Max Static Addresses Allowed	The number of source MAC addresses that can be manually added to the port security MAC address table for an interface. If the port link goes down, the statically configured MAC addresses remain in the MAC address table. The maximum number includes all dynamically-learned MAC addresses that have been converted to static MAC addresses.
Sticky Mode	The sticky MAC address learning mode, which is one of the following: <ul style="list-style-type: none"> • Enabled: MAC addresses learned or manually configured on this interface are learned in sticky mode. A sticky-mode MAC address is a MAC address that does not age out and is added to the running configuration. If the running configuration is saved to the startup configuration, the sticky addresses are saved to persistent storage and do not need to be relearned when the device restarts. Upon enabling sticky mode on an interface, all dynamically learned MAC addresses in the MAC address table for that interface are converted to sticky mode. Additionally, new addresses dynamically learned on the interface will also become sticky. • Disabled: When a link goes down on a port, all of the dynamically learned addresses are cleared from the source MAC address table the feature maintains. When the link is restored, the interface can once again learn addresses up to the specified limit. If sticky mode is disabled after being enabled on an interface, the sticky-mode addresses learned or manually configured on the interface are converted to dynamic entries and are automatically removed from persistent storage.
Violation Trap Mode	Whether the port security feature sends a trap to the SNMP agent when a port is locked and a frame with a MAC address not currently in the table arrives on the port. A port is considered to be locked once it has reached the maximum number of allowed dynamic or static MAC address entries in the port security MAC address table.
Violation Shutdown Mode	Whether the port security feature shuts down the port after MAC limit is reached.
Last Violation MAC/VLAN	The source MAC address and, if applicable, associated VLAN ID of the last frame that was discarded at a locked port.

VLAN MAC Locking

Use the **VLAN MAC Locking Status** page to configure VLAN MAC Locking. VLAN MAC locking allows you to secure the network by locking down allowable MAC addresses on a given VLAN. Packets with a matching source MAC address can be forwarded normally. All other packets will be discarded. VLAN MAC locking will lock the dynamic MAC entries.

If VLAN and port MAC locking are enabled, VLAN MAC locking will be given precedence over port MAC locking.

To access this page, click **Switching > Port Security > VLAN** in the navigation menu.

Table 242: Port Security Interface Configuration Fields

Field	Description
VLAN ID	The VLAN ID specified in the Ethernet frame received by the interface.
Interface	The interface associated with the rest of the data in the row.
Dynamic MAC Address	The MAC address that was learned on the device. An address is dynamically learned when a frame arrives on the interface and the source MAC address in the frame is added to the MAC address table.
Max Dynamic Addresses Allowed	The number of source MAC addresses that can be dynamically learned on an interface. If an interface reaches the configured limit, any other addresses beyond that limit are not learned, and the frames are discarded. Frames with a source MAC address that has already been learned will be forwarded. A dynamically-learned MAC address is removed from the MAC address table if the entry ages out, the link goes down, or the system reboots. Note that the behavior of a dynamically-learned address changes if the sticky mode for the interface is enabled or the address is converted to a static MAC address.
Operational MAC Limit	The number of source MAC addresses that are dynamically currently reached to that of Maximum Configured MAC Limit.
Violation Shutdown Mode	After MAC limit has reached, action will shut down the ports participating in the VLAN.
Violation Trap Mode	After MAC limit has reached, a log message will be generated with violation MAC address details.

To configure The VLAN MAC Locking, use the following buttons to perform the tasks:

- Use **Submit** to enable or disable VLAN MAC Locking Admin Mode.
- Use **Add** to configure VLAN MAC Locking.
- Use **Edit** to modify configuration parameters of VLAN MAC Locking.
- Use **Remove** to remove configured VLANs.

Port Security Statically Configured MAC Addresses

Use the **Port Security Static MAC Addresses** page to view static MAC addresses configured on an interface. From this page, you can delete statically configured MAC addresses.

To access this page, click **Switching > Port Security > Static MAC** in the navigation menu.

Table 243: Port Security Statically Configured MAC Address Fields

Field	Description
Interface	Select the physical interface or the <u>LAG</u> on which to view the dynamically learned MAC addresses.
MAC Address	This column lists the static MAC addresses, if any, configured on the selected port.
VLAN ID	Displays the VLAN ID corresponding to the statically configured MAC address.
Delete a static MAC Address	Enter the address of the statically configured MAC address to delete. All MAC addresses that are available to be deleted appear in the MAC Address - VLAN ID table.
VLAN ID	Enter the VLAN ID that corresponds to the statically configured MAC address to delete.

After you enter the MAC address and VLAN ID of the statically configured MAC address to delete, click **Submit** to remove the MAC address from the port and apply the new settings to the system. The screen refreshes, and the MAC address no longer appears in the table on the page.

Port Security Dynamically Learned MAC Addresses

Use the **Port Security Dynamic MAC Addresses** page to view a table with the dynamically learned MAC addresses on an interface. With dynamic locking, MAC addresses are learned on a “first arrival” basis. You specify how many addresses can be learned on the locked port.

To access this page, click **Switching > Port Security > Dynamic MAC** in the navigation menu.

Table 244: Port Security Dynamic Fields

Field	Description
Interface	Select the physical interface or the <u>LAG</u> on which to view the dynamically learned MAC addresses.
MAC Address	This column lists the dynamically learned MAC addresses, if any, on the selected port.
VLAN ID	Displays the VLAN ID corresponding to the dynamically learned MAC address.

Managing LLDP

The IEEE 802.1AB defined standard, LLDP (Link Layer Discovery Protocol), allows stations residing on an 802 LAN to advertise major capabilities and physical descriptions. This information is viewed by a network manager to identify system topology and detect bad configurations on the LAN.

LLDP is a one-way protocol; there are no request/response sequences. Information is advertised by stations implementing the transmit function, and is received and processed by stations implementing the receive function. The transmit and receive functions can be enabled/disabled separately per port. By default, both transmit and receive are disabled on all ports. The application is responsible for starting each transmit and receive state machine appropriately, based on the configured status and operational state of the port.

200 Series allows LLDP to have multiple LLDP neighbors per interface. The number of such neighbors is limited by the memory constraints. A product-specific constant defines the maximum number of

neighbors supported by the switch. There is no restriction on the number of neighbors supported on a per LLDP port. If all the remote entries on the switch are filled up, the new neighbors are ignored. In case of multiple VOIP devices on a single interface, the 802.1ab component sends the Voice VLAN configuration to all the VoIP devices.

Global Configuration

Use the [LLDP Global Configuration](#) page to specify LLDP parameters that are applied to the switch.

To access this page, click **Switching > LLDP > Global** in the navigation menu.

Table 245: LLDP Global Configuration Fields

Field	Description
Transmit Interval	Specifies the interval at which LLDP frames are transmitted. The default is 30 seconds, and the valid range is 1-32768 seconds.
Transmit Hold Multiplier	Specifies multiplier on the transmit interval to assign to TTL. The default is 4, and the range is 2-10.
Re-Initialization Delay	Specifies the delay before a re-initialization. The default is 2 seconds, and the range is 1-10 seconds.
Notification Interval	Limits the transmission of notifications. The default is 5 seconds, and the range is 5-3600 seconds.

If you make any changes to the page, click **Submit** to apply the new settings to the system.

LLDP Interface Configuration

Use the [LLDP Interface Configuration](#) page to specify [LLDP](#) parameters that are applied to a specific interface.

To access this page, click **Switching > LLDP > Interface** in the navigation menu.



Note

When adding or editing LLDP settings on an interface, select the appropriate checkbox to enable a feature, or clear the checkbox to disable a feature.

Table 246: LLDP Interface Summary Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. Only interfaces that have at least one LLDP setting enabled appear in the table. In the Add LLDP Interface window, use this field to select the interface with the LLDP settings to configure. In the Edit LLDP Interface window, this field identifies the interface that is being configured.
Port ID Subtype	The LLDP Port ID subtype of the interface, which is either MAC Address or Interface Name.
Link Status	The link status of the interface, which is either Up or Down. An interface that is down does not forward traffic.

Table 246: LLDP Interface Summary Fields (continued)

Field	Description
Transmit	The LLDP advertise (transmit) mode on the interface. If the transmit mode is enabled, the interface sends LLDP Data Units (LLDPDUs) that advertise the mandatory TLVs and any optional TLVs that are enabled.
Receive	The LLDP receive mode on the interface. If the receive mode is enabled, the device can receive LLDPDUs from other devices.
Notify	The LLDP remote data change notification status on the interface. If the notify mode is enabled, the interface sends <i>SNMP (Simple Network Management Protocol)</i> notifications when a link partner device is added or removed.
Optional TLV(s)	Select each checkbox next to the <i>TLV (Type Length Value)</i> information to transmit. Choices include: <ul style="list-style-type: none"> • System Name: To include system name TLV in LLDP frames. To configure the System Name, see System Description on page 39. • System Description: To include system description TLV in LLDP frames. • System Capabilities: To include system capability TLV in LLDP frames. • Port Description: To include port description TLV in LLDP frames. To configure the Port Description, see Port Description on page 127.
Transmit Management Information	Select the checkbox to enable the transmission of management address instance. Clear the checkbox to disable management information transmission. The default is disabled.

Use the buttons to perform the following tasks:

- To configure LLDP settings on an interface that does not have any LLDP settings enabled, click **Add**.
- To change the LLDP settings for an interface in the table, select the entry to update and click **Edit**. If you clear (disable) all LLDP settings, the entry is removed from the table.
- To clear (disable) all LLDP settings from one or more interfaces, select each entry to clear and click **Remove**.

After you click **Add** or **Edit**, a window opens and allows you to configure the LLDP settings for an interface. The following information describes the additional fields that appear in the windows used for adding or editing per-interface LLDP settings.

In addition to some of the fields that [Table 246](#) on page 241 describes, [Table 247](#) shows the additional fields available on the **Add LLDP Interface** window.

Table 247: LLDP Interface Add Fields

Field	Description
System Name	Select this option to include the user-configured system name in the LLDPDU the interface transmits. The system name is configured on the System Description page and is the SNMP server name for the device.
System Description	Select this option to include a description of the device in the LLDPDU the interface transmits. The description includes information about the product model and platform.

Table 247: LLDP Interface Add Fields (continued)

Field	Description
System Capabilities	Select this option to advertise the primary function(s) of the device in the LLDPDU the interface transmits.
Port Description	Select this option to include the user-configured port description in the LLDPDU the interface transmits.

If you make any changes to the page, click **Submit** to apply the new settings to the system.

Local Devices

Use the **LLDP Local Device Summary** page to view information about all interfaces on the device that are enabled to transmit [LLDP](#) information.

To access this page, click **Switching > LLDP > Local Devices** in the navigation menu.

Table 248: LLDP Local Devices Columns

Field	Description
Interface	The interface associated with the rest of the LLDP - 802.1AB data in the row. When viewing the details for an interface, this field identifies the interface that is being viewed.
Port ID	The port identifier, which is the physical address associated with the interface.
Port Description	A description of the port. An administrator can configure this information on the Port Description page.

Click **Refresh** to update the information on the screen with the most current data.

After you click **Details**, a window opens and displays additional information about the data the interface transmits in its LLDPDUs. The following information describes the additional fields that appear in the LLDP Local Device Information window.

Table 249: LLDP Local Devices Details

Field	Description
Chassis ID Subtype	The type of information used to identify the device in the Chassis ID field.
Chassis ID	The hardware platform identifier for the device.
Port ID Subtype	The type of information used to identify the interface in the Port ID field.
System Name	The user-configured system name for the device. The system name is configured on the System Description page and is the SNMP server name for the device.
System Description	The device description, which includes information about the product model and platform.
System Capabilities Supported	The primary function(s) the device supports.
System Capabilities Enabled	The primary function(s) the device supports that are enabled.

Table 249: LLDP Local Devices Details (continued)

Field	Description
Management Address	The physical address associated with the management interface of the device.
Management Address Type	The protocol type or standard associated with the management address.

Remote Devices

Use the **LLDP Remote Device Summary** page to view information about all interfaces on the device that are enabled to transmit LLDP information.

To access this page, click **Switching > LLDP > Remote Devices** in the navigation menu.

Table 250: LLDP Remote Device Summary Columns

Field	Description
Interface	The local interface that is enabled to receive LLDPDUs from remote devices.
Remote ID	The client identifier assigned to the remote system that sent the LLDPDU.
Chassis ID	The information the remote device sent as the Chassis ID TVL. This identifies the hardware platform for the remote system.
Port ID	The port on the remote system that transmitted the LLDP data.
System Name	The system name configured on the remote device.

Click **Refresh** to update the information on the screen with the most current data.

After you click **Details**, a window opens and displays additional information. If the interface has received LLDP data from a remote device, the window displays detailed information about the device. If the interface has not received any LLDPDUs from remote devices, the window displays a message indicating that no LLDP data has been received. The following information describes the additional fields that appear in the LLDP Remote Device Information window when LLDP data has been received on the selected interface.

Table 251: LLDP Remote Device Summary Columns

Field	Description
Chassis ID Subtype	The type of information used to identify the device in the Chassis ID field.
Port ID Subtype	The type of information used to identify the interface in the Port ID field.
System Description	The device description, which includes information about the product model and platform.
Port Description	The description of the port on the remote device that transmitted the LLDP data.
System Capabilities Supported	The primary function(s) the remote system supports. The possible capabilities include Other, Repeater, Bridge, <u>WLAN (Wireless Local Area Network)</u> AP, Router, Telephone, DOCSIS cable device, and Station.

Table 251: LLDP Remote Device Summary Columns (continued)

Field	Description
System Capabilities Enabled	The primary function(s) of the remote system that are both supported and enabled. The possible capabilities include Other, Repeater, Bridge, WLAN AP, Router, Telephone, DOCSIS cable device, and Station.
Time To Live	The number of seconds the local device should consider the LLDP data it received from the remote system to be valid.

Statistics

Use the **LLDP Statistics** page to view the global and interface *LLDP* statistics.

To access this page, click **Switching > LLDP > Statistics** in the navigation menu.

Table 252: LLDP Statistics Fields

Field	Description
System-wide Statistics	
Last Update	Displays the time when an entry was created, modified, or deleted in the tables associated with the remote systems.
Total Inserts	Displays the number of times a complete set of information advertised by a particular MAC Service Access Point (MSAP) has been inserted into the tables associated with the remote systems.
Total Deletes	Displays the number of times a complete set of information advertised by a particular MAC Service Access Point (MSAP) has been deleted from the tables associated with the remote systems.
Total Drops	Displays the number of times a complete set of information advertised by a particular MAC Service Access Point (MSAP) could not be entered into tables associated with the remote systems because of insufficient resources.
Total Ageouts	Displays the number of times a complete set of information advertised by a particular MAC Service Access Point (MSAP) has been deleted from tables associated with the remote systems because the information timelines interval has expired.
Port Statistics	
Interface	Identifies the interfaces.
Transmit Total	Displays the total number of LLDP frames transmitted by the LLDP agent on the corresponding port.
Receive Total	Displays the total number of valid LLDP frames received by the LLDP agent on the corresponding port, while the LLDP agent is enabled.
Discards	Displays the number of LLDP TLVs discarded for any reason by the LLDP agent on the corresponding port.
Errors	Displays the number of invalid LLDP frames received by the LLDP agent on the corresponding port, while the LLDP agent is enabled.

Table 252: LLDP Statistics Fields (continued)

Field	Description
Ageouts	Displays the number of age-outs that occurred on a given port. An age-out is the number of times the complete set of information advertised by a particular MAC Service Access Point (MSAP) has been deleted from tables associated with remote entries because the information timeliness interval had expired.
TLV Discards	Displays the number of LLDP TLVs (Type, Length, Value sets) discarded for any reason by the LLDP agent on the corresponding port.
TLV Unknowns	Displays the number of LLDP TLVs received on the local ports which were not recognized by the LLDP agent on the corresponding port.
TLV MED	Displays the total number of LLDP-MED TLVs received on the local ports.
TLV 802.1	Displays the total number of LLDP TLVs received on the local ports which are of type 802.1.
TLV 802.3	Displays the total number of LLDP TLVs received on the local ports which are of type 802.3.

- Click **Refresh** to update the page with the most current information.
- Click **Clear** to clear the LLDP statistics of all the interfaces.

LLDP-MED

The Link Layer Discovery Protocol-Media Endpoint Discovery (LLDP-MED) is an enhancement to [LLDP](#) that features:

- Auto-discovery of LAN policies (such as VLAN, Layer 2 Priority and DiffServ settings), enabling plug and play networking.
- Device location discovery for creation of location databases.
- Extended and automated power management of Power over Ethernet endpoints.
- Inventory management, enabling network administrators to track their network devices and determine their characteristics (manufacturer, software and hardware versions, serial/asset number).

LLDP-MED Global Configuration

Use the **LLDP-MED Global Configuration** page to set global parameters for LLDP-MED operation.

To access this page, click **Switching > LLDP-MED > Global** in the navigation menu.

Table 253: LLDP Global Configuration Fields

Field	Description
Fast Start Repeat Count	Specifies the number of LLDP PDUs that will be transmitted when the protocol is enabled. The range is from (1 to 10). The default value s 3.
Device Class	<p>Specifies local device's MED Classification. The following three represent the actual endpoints:</p> <ul style="list-style-type: none"> • Class I Generic [IP Communication Controller etc.] • Class II Media [Conference Bridge etc.] • Class III Communication [IP Telephone etc.] <p>The fourth device is Network Connectivity Device, which is typically a LAN switch/router, IEEE 802.1 bridge, IEEE 802.11 wireless access point, etc.</p>

Click **Submit** to updated the switch. The changes take effect but will not be retained across a power cycle unless a save is performed.

LLDP-MED Interface Configuration

Use the **LLDP-MED Interface Configuration** page to enable LLDP-MED mode on an interface and to configure its properties. To configure the settings for one or more interfaces, select each entry to modify and click **Edit**. The same LLDP-MED settings are applied to all selected interfaces.

To access this page, click **Switching > LLDP-MED > Interface** in the navigation menu.

Table 254: LLDP-MED Interface Configuration Fields

Field	Description
Interface	Selects the port that you want to configure LLDP-MED — 802.1AB on. You can select All to configure all interfaces on the DUT with the same properties. The Interface Configuration page will not be able to display the summary of 'All' interfaces. The summary of individual interfaces is visible from the Interface Configuration page. The Interface Configuration page for the 'All' option will always display the LLDP-MED mode and notification mode as 'disabled' and checkboxes for 'Transmit TLVs' will always be unchecked.
Link Status	The link status of the interface, which is either Up or Down. An interface that is down does not forward traffic.
MED Status/LLDP-MED Mode	The administrative status of LLDP-MED on the interface. When LLDP-MED is enabled, the transmit and receive function of LLDP is effectively enabled on the interface.
Notification Status/Configuration Notification Mode	Whether LLDP-MED topology change notifications are enabled or disabled on the interface.
Operational Status	Whether the interface will transmit TLVs.
Transmit TLVs	<p>The LLDP-MED TLV(s) that the interface transmits:</p> <ul style="list-style-type: none"> • MED Capabilities: 0 • Network Policy: 1

Click **Submit** to send the updated configuration to the switch. These changes take effect immediately but will not be retained across a power cycle unless a save is performed.

LLDP Local Device Information

The **LLDP-MED Local Device Information** page displays information on LLDP-MED information advertised on the selected local interface. To access this page, click **Switching > LLDP-MED > Local Devices** in the navigation menu.

Table 255: LLDP-MED Local Device Information Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When viewing LLDP-MED details for an interface, this field identifies the interface that is being viewed.
Port ID	The MAC address of the interface. This is the MAC address that is advertised in LLDP-MED PDUs. After you click Details , a window opens and shows detailed information about the LLDP-MED information the selected interface transmits. The following information describes the additional fields that appear in the LLDP-MED Local Device Information window.
Network Policy Information The information in this table identifies the data transmitted in the Network Policy TLVs.	
Media Application Type	The media application type transmitted in the TLV. The application types are unknown, voicesignaling, guestvoice, guestvoicesignalling, softphonevoice, videoconferencing, streamingvideo, videosignalling. Each application type that is transmitted has the VLAN ID, priority, DSCP, tagged bit status and unknown bit status. A port may transmit one or many such application types. This information is displayed only when a network policy TLV has been transmitted.
VLAN ID	The VLAN ID associated with a particular policy type.
Priority	The user priority associated with a particular policy type.
DSCP	The DSCP value associated with a particular policy type.
Unknown Bit Status	The unknown bit associated with a particular policy type.
Tagged Bit Status	Identifies whether the network policy is defined for tagged or untagged VLANs.
Location Information	
Sub Type	The type of location information: <ul style="list-style-type: none"> • Coordinate Based: The location map coordinates (latitude, longitude and altitude) of the device. • Civic Address: The civic or street address location of the device. • ELIN: The Emergency Call Service (ECS) Emergency Location Identification Number (ELIN) of the device.
Information	This column displays the information related to the coordinates, civic address, and ELIN for the device.

Click **Refresh** to update the page with the latest information from the router.

LLDP-MED Remote Device Information

The **LLDP-MED Remote Device Information** page displays information about the remote devices the local system has learned about through the LLDP-MED data units received on its interfaces. Information is available about remote devices only if an interface receives an LLDP-MED data unit from a device. To view additional information about a remote device, select the interface that received the LLDP-MED data and click **Details**. The information below is organized according to the order in which the fields appear in the LLDP-MED Remote Device Information window.

To access this page, click **Switching > LLDP-MED > Remote Devices** in the navigation menu.

Table 256: LLDP-MED Remote Device Information Fields

Field	Description
Interface	The local interface that has received LLDP-MED data units from remote devices.
Remote ID	The client identifier assigned to the remote system that sent the LLDP-MED data unit.
Capability Information	
Supported Capabilities	The supported capabilities that were received in the MED TLV on this interface.
Enabled Capabilities	The supported capabilities on the remote device that are also enabled.
Device Class	<p>The MED Classification advertised by the TLV from the remote device. The following three classifications represent the actual endpoints:</p> <ul style="list-style-type: none"> • Class I Generic (for example, IP Communication Controller) • Class II Media (for example, Conference Bridge) • Class III Communication (for example, IP Telephone) <p>The fourth device is Network Connectivity Device, which is typically a device such as a LAN switch or router, IEEE 802.1 bridge, or IEEE 802.11 wireless access point.</p>
Network Policy Information This section describes the information in the network policy TLVs received in the LLDP-MED frames on this interface.	
Media Application Type	The media application type received in the TLV from the remote device. The application types are unknown, voicesignaling, guestvoice, guestvoicesignalling, softphonevoice, videoconferencing, streamingvideo, videosignalling. Each application type that is transmitted has the VLAN ID, priority, DSCP, tagged bit status and unknown bit status. The port on the remote device may transmit one or many such application types. This information is displayed only when a network policy TLV has been received.
VLAN ID	The VLAN ID associated with a particular policy type.
Priority	The user priority associated with a particular policy type.
DSCP	The DSCP value associated with a particular policy type.
Unknown Bit Status	The unknown bit associated with a particular policy type.
Tagged Bit Status	Identifies whether the network policy is defined for tagged or untagged VLANs.
Inventory Information This section describes the information in the inventory TLVs received in the LLDP-MED frames on this interface.	
Hardware Revision	The hardware version advertised by the remote device.
Firmware Revision	The firmware version advertised by the remote device.

Table 256: LLDP-MED Remote Device Information Fields (continued)

Field	Description
Software Revision	The software version advertised by the remote device.
Serial Number	The serial number advertised by the remote device.
Manufacturer Name	The name of the system manufacturer advertised by the remote device.
Model Name	The name of the system model advertised by the remote device.
Asset ID	The system asset ID advertised by the remote device.
Location Information This section describes the information in the location TLVs received in the LLDP-MED frames on this interface.	
Sub Type	The type of location information advertised by the remote device.
Information	The text description of the location information included in the subtype.
Extended <i>PoE (Power over Ethernet)</i>	Whether the remote device is advertised as a PoE device.
Device Type	If the remote device is a PoE device, this field identifies the PoE device type of the remote device connected to this port.

Click **Refresh** to update the page with the latest information from the router.

Loop Protection

L2 Loop Protection feature allows loop detection in downstream switches that do not run spanning tree. It can optionally disable the associated port on loop detection.

The Loop Protection feature is not intended for ports that serve as uplinks between spanning tree aware switches. Loop Protection feature is designed for unmanaged switches which drop spanning Tree BPDUs. This feature detects physical and logical loops between Ethernet ports on a device. The feature needs to be enabled globally before enabling it at the interface level for the system policy filter to be installed.

Loop Protection Configuration

Use the **Loop Protection Configuration** page to configure the Loop Protection feature. Loops on a network consume resources and can impact network performance. When loop protection is enabled on the switch and on one or more interfaces (ports and trunks), the interfaces send loop protection protocol data units (PDUs) to the multicast destination address 01:80:C2:00:00:08. When an interface receives a loop protection PDU, it compares the source MAC address with its own. If the MAC addresses match, a loop is detected and a configured action is taken, which may include shutting down the port for a specified period. An interface can also be configured to receive and take action in response to loop protection PDUs, but not to send out the PDUs itself.

To access this page, click **Switching > Loop Protection > Configuration** in the navigation menu.

Table 257: Loop Protection Configuration Fields

Field	Description
Loop Protection	Enables or disables the loop protection feature globally on the switch. The loop protection feature is not supported on dynamic trunks. The loop protection feature will be automatically disabled if it was previously enabled on a static trunk that is now configured as dynamic.
Transmission Time (Seconds)	The interval at which the switch sends loop protection PDUs on interfaces that are enabled to send them.
Maximum PDU Received	The count of loop protection packets received by the switch after which the interface will be err-disabled.
Interface	The port or trunk ID. Select an interface and click Edit to edit the Loop Protection port configuration. Click Edit All to apply the same configuration to all interfaces.
Action	The action to be taken when a loop is detected on the port: <ul style="list-style-type: none"> • Shutdown Port: Shut down the port for the configured Transmission Time. • Shutdown Port and Log: Shut down the port for the configured Transmission Time and send a message to the system log. • Log Only: Send a message to the system log but do not shut down the port.
Status	The current status of the interface. Link Up indicates the interface is operating normally. Link Down indicates that the port has been shut down due to the detection of a loop.
Loop	Whether a loop is currently detected on the interface. If blank, then no loop is detected.
Loop Count	The number of times a loop has occurred on the interface.
Time of Last Loop	The date and time the most recent loop was detected.

Click **Submit** to updated the switch. The changes take effect but will not be retained across a power cycle unless a save is performed.

Multiple Registration Protocol Configuration

Like 802.1AS, Multiple Registration Protocol (MRP) is an Audio Video Bridging (AVB) feature that is available on some 200 Series platforms. MVR is a base registration protocol that enables devices running an MRP application to register attributes to other devices in a network. MRP provides an application to register attributes such as bandwidth requirement for a given AV stream and MAC address information. It is used by various applications to propagate the registration. 200 Series switches support the following MRP applications:

- Multiple MAC Registration Protocol (MMRP)
- Multiple Stream Reservation Protocol (MSRP)
- Multiple VLAN Registration Protocol (MVRP)

MMRP allows for the propagation MAC address information in the network, and allows for the registration and deregistration of both individual MAC address information and group MAC address membership. End stations may request to join or leave a multicast group, or to register an individual MAC address with a specific VLAN. MAC address entries can be dynamically registered and deregistered if MMRP is administratively enabled on the switch.

MSRP reserves necessary resources in the network to facilitate time sensitive traffic to flow end to end. In a typical network, there are multiple Talkers (those who transmit streams) and multiple Listeners (those who receive streams from one or many Talkers). Each flow has specific bandwidth, frame rate, and time sync requirements. With the use of MSRP these resources are guaranteed through all intermediate devices that are between any talker and listener.

MVRP registers VLANs in the network, enabling automatic VLAN configuration on the switch. In a typical network, VLAN tagging is common. Many nodes require ingress traffic to be tagged with specific VLAN ID, and other nodes require egress traffic to be transmitted with a specific VLAN ID. With the use of MVRP on both ingress and egress, no manual VLAN configuration is required to pass tagged traffic through the network.

**Note**

MRP framework must be available and enabled in all intermediate devices to ensure that the propagation of the attributes occurs throughout the network.

With MRP, network attributes are declared, registered, withdrawn, and removed completely dynamically without any user intervention. This dynamic nature is especially useful in networks where:

- Network attributes are likely to change frequently, requiring reconfiguration of the intermediate devices.
- Recipients of these attributes frequently increase or decrease in number.
- Each of these changes without a dynamic self-adjusting framework would require constant attention from the network administrator.

MRP Configuration

Use the **MRP Configuration** page to configure global MRP settings for the switch. To access this page click **Switching > MRP>Configuration**.

**Note**

The fields available on the MRP Configuration page vary based on the platform and its supported features.

Table 258: MRP Configuration Fields

Field	Description
MVRP	<p>Multiple VLAN Registration Protocol (MVRP) registers VLANs in the network, enabling automatic VLAN configuration on the device. In a typical network, VLAN tagging is common. Many nodes require ingress traffic to be tagged with a specific VLAN ID, and other nodes require egress traffic to be transmitted with a specific VLAN ID. With the use of MVRP on both ingress and egress, no manual VLAN configuration is required to pass tagged traffic through the network.</p> <p>Admin Mode: The administrative mode of MVRP on the device.</p> <p>Periodic State Machine: Select this option to enable the MRP Periodic State Machine for MVRP on the device. When enabled, the state machine can help limit the effect of topology changes and reduce the number of protocol data units (PDUs) transmitted between devices.</p>
MMRP	<p>Multiple MAC Registration Protocol (MMRP) allows the propagation of MAC address information in the network, and allows for the registration and deregistration of both individual MAC address information and group MAC address membership. End stations may request to join or leave a multicast group, or to register an individual MAC address with a specific VLAN. MAC address entries can be dynamically registered and deregistered if MMRP is administratively enabled on the device.</p> <p>Admin Mode: The administrative mode of MMRP on the device.</p> <p>Periodic State Machine: Select this option to enable the MRP Periodic State Machine for MMRP on the device. When enabled, the state machine can help limit the effect of topology changes and reduce the number of protocol data units (PDUs) transmitted between devices.</p>
MSRP	<p>Multiple Stream Reservation Protocol (MSRP) reserves necessary resources in the network to facilitate the end-to-end flow of time sensitive traffic. In a typical network, there are multiple Talkers (those who transmit streams) and multiple Listeners (those who receive streams from one or many Talkers). Each flow has specific bandwidth, frame rate, and time sync requirements. With the use of MSRP these resources are guaranteed through all intermediate devices that are between any Talker and Listener.</p> <p>Admin Mode: The administrative mode of MSRP on the device.</p> <p>Periodic State Machine: Select this option to enable the MRP Periodic State Machine for MSRP on the device. When enabled, the state machine can help limit the effect of topology changes and reduce the number of protocol data units (PDUs) transmitted between devices.</p>
Talker Pruning	Select this option to enable MSRP Talker pruning. The MSRP Talker is the source of an AV stream. When enabled, Talker pruning stops MSRP declarations sent by the talker unless a Listener registers and requests them.
Boundary Propagation	Select this option to enable MSRP boundary propagation on the device.
Max Fan In Ports	The maximum number of ports where MSRP registrations are allowed.

MRP Interface Configuration

Use the **MRP Interface Configuration** page to view and configure the per-interface Multiple Registration Protocol (MRP) settings. To change the current settings for one or more interfaces, select each interface to modify and click **Edit**. The same MRP settings are applied to all selected interfaces.

To access this page, click **Switching > MRP > Interface** in the navigation menu. In the following image, the MMRP mode on ports g4 and g5 is being enabled.

To configure one or more ports or LAGs, select the checkbox next to each port or LAG to configure. You can select multiple ports to apply the same settings to the selected interfaces.

Table 259: MRP Port Configuration Fields

Field	Description
Interface	Identifies the interface associated with the rest of the information in the row.
MVRP Mode	The administrative mode of Multiple VLAN Registration Protocol (MVRP) on the interface. MVRP registers VLANs in the network, enabling automatic VLAN configuration on the device.
MMRP Mode	The administrative mode of Multiple MAC Registration Protocol (MMRP) on the interface. MMRP allows the propagation of MAC address information in the network and allows for the registration and deregistration of both individual MAC address information and group MAC address membership.
MSRP Mode	The administrative mode of Multiple Stream Reservation Protocol (MSRP) on the interface. MSRP reserves necessary resources in the network to facilitate the end-to-end flow of time sensitive traffic.
MSRP SR Class PVID	The default VLAN ID to be used for MSRP stream traffic.
Join Timer	The amount of time to wait for JoinIn messages from other MRV participants after the interface sends a Join message. If the amount of time specified in this field passes before the interface receives a JoinIn message, the interface resends the Join message.
Leave Timer	The amount of time to wait before the interface deregisters attributes from other MRV participants. If the interface receives Join messages from other participants before the Leave timer expires, the attributes are not deregistered.
Leave All Timer	The amount of time to wait, after the interface starts the MRP registration process, before the participants refresh and re-register their attributes.

MMRP Statistics

The **MMRP Statistics** page displays information regarding the MMRP frames transmitted and received by the switch and by each interface. To access this page click **Switching > MRP > Advanced > MMRP Statistics**.

Table 260: MMRP Statistics Fields

Field	Description
Interface	In the Interface Statistics table, this field identifies the interface associated with the rest of the data in the row.
Frames Received	Shows the number of MMRP frames which were received on the switch.
Bad Header	Shows number of MMRP frames with bad headers which were received on the switch.
Bad Format	Shows number of MMRP frames with bad PDUs body formats which were received on the switch.

Table 260: MMRP Statistics Fields (continued)

Field	Description
Frames Transmitted	Shows number of MMRP frames which were transmitted on the switch.
Transmission Failures	Shows number of MMRP frames that the switch failed to transmit.

To reload the page, click **Refresh**. To clear the statistics for one or more ports, select the checkbox next to the interface or interfaces, and click **Clear**. To clear the statistics for all interfaces, select the checkbox in the heading row, and click **Clear**.

MVRP Statistics

The **MVRP Statistics** page displays information about the MVRP frames transmitted and received by the switch and by each interface. To access this page, click **Switching > MRP>MVRP Statistics**.

Table 261: MVRP Statistics

Field	Description
Interface	In the Interface Statistics table, this field identifies the interface associated with the rest of the data in the row.
Frames Received	Shows number of MVRP frames that have been received on the switch.
Bad Header	Shows number of MVRP frames with bad headers that have been received on the switch.
Bad Format	Shows number of MVRP frames with bad PDUs body formats that have been received on the switch.
Frames Transmitted	Shows number of MVRP frames which that have been transmitted on the switch.
Transmission Failures	Shows number of MVRP frames the switch failed to transmit.
Message Failures	Shows the number of messages that failed to be added to the queue.
Registration Failures	Shows the number of MVRP frames that failed to register on a device or particular interface.

To reload the page, click **Refresh**. To clear the statistics for one or more ports, select the checkbox next to the interface or interfaces, and click **Clear**. To clear the statistics for all interfaces, select the checkbox in the heading row, and click **Clear Counters**.

5 Configuring Routing

Configuring ARP
Configuring Global IP Settings
Router
Configuring Routing Information Protocol (RIP)

200 Series software supports IP routing. Use the links in the **Routing** navigation menu to manage routing on the system.

When a packet enters the switch, the destination MAC address is checked to see if it matches any of the configured routing interfaces. If it does, then the silicon searches the host table for a matching destination IP address. If an entry is found, then the packet is routed to the host. If there is not a matching entry, then the switch performs a longest prefix match on the destination IP address. If an entry is found, then the packet is routed to the next hop. If there is no match, then the packet is routed to the next hop specified in the default route. If there is no default route configured, then the packet is passed to the 6200 series software to be handled appropriately.

The routing table can have entries added either statically by the administrator or dynamically via a routing protocol. The host table can have entries added either statically by the administrator or dynamically via ARP.



Note

200 Series supports the [*BGP \(Border Gateway Protocol\)*](#). BGP is available as a separate module and might not be available on all platforms. The BGP features can be configured only by using the CLI. No web-based administrative pages are available for BGP configuration.

Configuring ARP

The [*ARP \(Address Resolution Protocol\)*](#) protocol associates a Layer 2 MAC address with a Layer 3 IPv4 address. 200 Series software features both dynamic and manual ARP configuration. With manual ARP configuration, you can statically add entries into the ARP table.

ARP is a necessary part of the internet protocol (IP) and is used to translate an IP address to a media (MAC) address, defined by a local area network (LAN) such as Ethernet. A station needing to send an IP packet must learn the MAC address of the IP destination, or of the next hop router, if the destination is not on the same subnet. This is achieved by broadcasting an ARP request packet, to which the intended recipient responds by unicasting an ARP reply containing its MAC address. Once learned, the MAC address is used in the destination address field of the Layer 2 header prepended to the IP packet.

The ARP cache is a table maintained locally in each station on a network. ARP cache entries are learned by examining the source information in the ARP packet payload fields, regardless of whether it is an ARP request or response. Thus, when an ARP request is broadcast to all stations on a LAN segment or virtual LAN (VLAN), every recipient has the opportunity to store the sender's IP and MAC address in their respective ARP cache. The ARP response, being unicast, is normally seen only by the requestor,

who stores the sender information in its ARP cache. Newer information always replaces existing content in the ARP cache.

The number of supported ARP entries is platform-dependent.

Devices can be moved in a network, which means the IP address that was at one time associated with a certain MAC address is now found using a different MAC, or may have disappeared from the network altogether (that is, it has been reconfigured, disconnected, or powered off). This leads to stale information in the ARP cache unless entries are updated in reaction to new information seen on the network, periodically refreshed to determine if an address still exists, or removed from the cache if the entry has not been identified as a sender of an ARP packet during the course of an ageout interval, usually specified via configuration.

The **Routing > ARP Table** folder contains links to the following web pages that configure and display ARP detail:

- [ARP Create](#) on page 257
- [ARP Table Configuration](#) on page 258

ARP Create

Use the **ARP Create** page to add an entry to the Address Resolution Protocol table.

To access this page, click **Routing > ARP Table > Summary** in the navigation menu.

The ARP Table displays at the bottom of the page, and contains the following fields:

Use the buttons to perform the following tasks:

- To add a static ARP entry, click **Add**. The Add Static ARP Entry dialog box opens. Specify the new entry information in the available fields.
- To delete one or more ARP entries, select each entry to delete and click **Remove**. Note that ARP entries designated as Local cannot be removed.

Table 262: ARP Create Fields

Field	Description
IP Address	The IP address of a network host on a subnet attached to one of the device's routing interfaces. When adding a static ARP entry, specify the IP address for the entry after you click Add .
MAC Address	The unicast MAC address (hardware address) associated with the network host. When adding a static ARP entry, specify the MAC address to associate with the IP address in the entry.
Interface	The routing interface associated with the ARP entry. The network host is associated with the device through this interface.

Table 262: ARP Create Fields (continued)

Field	Description
Type	The ARP entry type: <ul style="list-style-type: none"> • Dynamic: An ARP entry that has been learned by the router • Gateway: A dynamic ARP entry that has the IP address of a routing interface • Local: An ARP entry associated with the MAC address of a routing interface on the device • Static: An ARP entry configured by the user
Age	The age of the entry since it was last learned or refreshed. This value is specified for Dynamic or Gateway entries only (it is left blank for all other entry types).

After you enter an IP address and the associated MAC address, click **Submit** to apply the changes to the system and create the entry in the ARP table.

ARP Table Configuration

Use the **ARP Table Configuration** page to change the configuration parameters for the Address Resolution Protocol Table. You can also use this screen to display the contents of the table.

To access this page, click **Routing > ARP Table > Configuration** in the navigation menu.

Table 263: ARP Table Configuration Fields

Field	Description
Age Time	The amount of time, in seconds, that a dynamic ARP entry remains in the ARP table before aging out.
Response Time	The amount of time, in seconds, that the device waits for an ARP response to an ARP request that it sends.
Retries	The maximum number of times an ARP request will be retried after an ARP response is not received. The number includes the initial ARP request.
Cache Size	The maximum number of entries allowed in the ARP table. This number includes all static and dynamic ARP entries.
Dynamic Renew	When selected, this option allows the ARP component to automatically attempt to renew dynamic ARP entries when they age out.

If you make any changes to the page, click **Submit** to apply the changes to the system.

Configuring Global IP Settings

The **Routing > IP** folder contains links to the following web pages that configure and display IP routing data:

- [Routing IP Configuration](#) on page 259
- [Interface Summary](#) on page 260
- [Routing IP Interface Configuration](#) on page 262

- [Routing IP Loopback Configuration](#) on page 264
- [Routing IP Statistics](#) on page 264

Routing IP Configuration

Use the **Routing IP Configuration** page to configure global routing settings on the device. Routing provides a means of transmitting IP packets between subnets on the network. Routing configuration is necessary only if the device is used as a Layer 3 device that routes packets between subnets. If the device is used as a Layer 2 device that handles switching only, it typically connects to an external Layer 3 device that handles the routing functions; therefore, routing configuration is not required on the Layer 2 device.

To access this page, click **Routing > IP > Configuration** in the navigation menu.

Table 264: Configuration Fields

Field	Description
Routing Mode	The administrative mode of routing on the device. The options are as follows: Enable – The device can act as a Layer 3 device by routing packets between interfaces configured for IP routing. Disable – The device acts as a Layer 2 bridge and switches traffic between interfaces. The device does not perform any internetwork routing.
ICMP Echo Replies	Select Enable or Disable from the drop-down menu. If you select Enable , then only the router can send ECHO replies. By default, <i>ICMP (Internet Control Message Protocol)</i> Echo Replies are sent for echo requests.
ICMP Redirects	Select this option to allow the device to send ICMP Redirect messages to hosts. An ICMP Redirect message notifies a host when a better route to a particular destination is available on the network segment.
ICMP Rate Limit Interval	To control the ICMP error packets, you can specify the number of ICMP error packets that are allowed per burst interval. By default, the rate limit is 100 packets per second; that is, the burst interval is 1000 milliseconds. To disable ICMP rate limiting, set this field to 0. The valid range is 0 to 2147483647 milliseconds.
ICMP Rate Limit Burst Size	To control the ICMP error packets, you can specify the number of ICMP error packets that are allowed per burst interval. By default, the burst size is 100 packets. When the burst interval is zero, then configuring this field is not a valid option. The valid burst size range is 1 to 200.
Static Route Preference	The default distance (preference) for static routes. Lower route-distance values are preferred when determining the best route. The value configured for Static Route Preference is used when using the CLI to configure a static route and no preference is specified. Changing the Static Route Preference does not update the preference of existing static routes.
Local Route Preference	The default distance (preference) for local routes.
Maximum Next Hops	The maximum number of hops supported by the switch. This is a read-only value.

Table 264: Configuration Fields (continued)

Field	Description
Maximum Routes	The maximum number of routes (routing table size) supported by the switch.
Global Default Gateway	<p>The IP address of the default gateway for the device. If the destination IP address in a packet does not match any routes in the routing table, the packet is sent to the default gateway. The gateway specified in this field is more preferred than a default gateway learned from a <i>DHCP (Dynamic Host Configuration Protocol)</i> server. Use the icons associated with this field to perform the following tasks:</p> <ul style="list-style-type: none"> To configure the default gateway, click the Edit icon and specify the IP address of the default gateway in the available field. To reset the IP address of the default gateway to the factory default value, click the Reset icon associated with this field.

If you make any changes to the page, click **Submit** to apply the changes to the system.

Interface Summary

The **Routing IP Interface Summary** shows summary information about the routing configuration for all interfaces. To view additional routing configuration information for an interface, select the interface with the settings to view and click **Details**.

To access this page, click **Routing > IP > Interface Summary** in the navigation menu.

Table 265: Interface Summary Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When viewing details about the routing settings for an interface, this field identifies the interface being viewed.
Status	Whether the interface is capable of routing IP packets (Up) or cannot route packets (Down). For the status to be Up, the routing mode and administrative mode for the interface must be enabled. Additionally, the interface must have an IP address and be physically up (active link).
IP Address	The IP address of the interface.
Subnet Mask	The IP subnet mask for the interface (also known as the network mask or netmask). It defines the portion of the interface's IP address that is used to identify the attached network.
Admin Mode	The administrative mode of the interface, which is either Enabled or Disabled.
State	The state of the interface, which is either Active or Inactive. An interface is considered active if the link is up, and the interface is in a forwarding state.
MAC Address	The burned-in physical address of the interface. The format is six two-digit hexadecimal numbers separated by colons, for example 00:06:29:32:81:40.

Table 265: Interface Summary Fields (continued)

Field	Description
Proxy ARP	Whether proxy ARP is enabled or disabled on the interface. When proxy ARP is enabled, the interface can respond to an ARP request for a host other than itself. An interface can act as an ARP proxy if it is aware of the destination and can route packets to the intended host, which is on a different subnet than the host that sent the ARP request.
IP MTU	The largest IP packet size the interface can transmit, in bytes. The IP Maximum Transmission Unit (MTU) is the maximum frame size minus the length of the Layer 2 header.

After you click **Details**, the **Details** window opens and displays detailed routing information for the selected interface. The following information describes the fields in this window that are not displayed on the summary page.

Table 266: Interface Summary Details Fields

Field	Description
Routing Mode	Whether routing is administratively enabled or disabled on the interface.
Link Speed Data Rate	The physical link data rate of the interface.
IP Address Configuration Method	The source of the IP address, which is one of the following: <ul style="list-style-type: none"> • None: The interface does not have an IP address. • Manual: The IP address has been statically configured by an administrator. • DHCP: The IP address has been learned dynamically through <i>DHCP</i>. If the method is DHCP but the interface does not have an IP address, the interface is unable to acquire an address from a network DHCP server.
Bandwidth	The configured bandwidth on this interface. This setting communicates the speed of the interface to higher-level protocols.
Encapsulation Type	The link layer encapsulation type for packets transmitted from the interface, which can be either Ethernet or SNAP.
Forward Net Directed Broadcasts	Indicates how the interface handles network-directed broadcast packets. A network-directed broadcast is a broadcast directed to a specific subnet. The possible values are as follows: <ul style="list-style-type: none"> • Enabled: Network directed broadcasts are forwarded. • Disabled: Network directed broadcasts are dropped.
Local Proxy ARP	Whether local proxy ARP is enabled or disabled on the interface. When local proxy ARP is enabled, the interface can respond to an ARP request for a host other than itself. Unlike proxy ARP, local proxy ARP allows the interface to respond to ARP requests for a host that is on the same subnet as the host that sent the ARP request. This feature is useful when a host is not permitted to reply to an ARP request from another host in the same subnet, for example when using the protected ports feature.

Table 266: Interface Summary Details Fields (continued)

Field	Description
Destination Unreachables	Whether the interface is allowed to send <i>ICMP</i> Destination Unreachable message to a host if the intended destination cannot be reached for some reason. If the status of this field is Disabled, this interface will not send ICMP Destination Unreachable messages to inform the host about the error in reaching the intended destination.
ICMP Redirects	Whether the interface is allowed to send ICMP Redirect messages. The device sends an ICMP Redirect message on an interface only if ICMP Redirects are enabled both globally and on the interface. An ICMP Redirect message notifies a host when a better route to a particular destination is available on the network segment.

Routing IP Interface Configuration

Use the [Routing IP Interface Configuration](#) page to configure the IP routing settings for each interface.

To access this page, click **Routing > IP > Interface Configuration** in the navigation menu.

Table 267: Interface Configuration Fields

Field	Description
Interface	The menu contains all interfaces that can be configured for routing. To configure routing settings for an interface, select it from the menu and then configure the rest of the settings on the page.
Status	Whether the interface is currently capable of routing IP packets (Up) or cannot route packets (Down). For the status to be Up, the routing mode and administrative mode for the interface must be enabled. Additionally, the interface must have an IP address and be physically up (active link).
Routing Mode	The administrative mode of IP routing on the interface.
Admin Mode	The administrative mode of the interface. If an interface is administratively disabled, it cannot forward traffic.
State	The state of the interface, which is either Active or Inactive. An interface is considered active if the link is up, and the interface is in a forwarding state.
Link Speed Data Rate	The physical link data rate of the interface.
IP Address Configuration Method	The method to use for configuring an IP address on the interface, which can be one of the following: <ul style="list-style-type: none"> • None: No address is to be configured. • Manual: The address is to be statically configured. When this option is selected you can specify the IP address and subnet mask in the available fields. • DHCP: The interface will attempt to acquire an IP address from a network <i>DHCP</i> server.

Table 267: Interface Configuration Fields (continued)

Field	Description
IP Address	The IP address of the interface. This field can be configured only when the selected IP Address Configuration Method is Manual. If the method is DHCP, the interface attempts to lease an IP address from a DHCP server on the network, and the IP address appears in this field (read-only) after it is acquired. If this field is blank, the IP Address Configuration Method might be None, or the method might be DHCP and the interface is unable to lease an address.
Subnet Mask	The IP subnet mask for the interface (also known as the network mask or netmask). This field can be configured only when the selected IP Address Configuration Method is Manual.
MAC Address	The burned-in physical address of the interface. The format is six two-digit hexadecimal numbers separated by colons, for example 00:06:29:32:81:40.
IP MTU	The largest IP packet size the interface can transmit, in bytes. The IP Maximum Transmission Unit (MTU) is the maximum frame size minus the length of the Layer 2 header.
Bandwidth	The configured bandwidth on this interface. This setting communicates the speed of the interface to higher-level protocols.
Encapsulation Type	The link layer encapsulation type for packets transmitted from the interface, which can be either Ethernet or SNAP.
Forward Net Directed Broadcasts	Determines how the interface handles network-directed broadcast packets. A network-directed broadcast is a broadcast directed to a specific subnet. If this option is selected, network directed broadcasts are forwarded. If this option is clear, network directed broadcasts are dropped.
Proxy ARP	When this option is selected, proxy ARP is enabled, and the interface can respond to an ARP request for a host other than itself. An interface can act as an ARP proxy if it is aware of the destination and can route packets to the intended host, which is on a different subnet than the host that sent the ARP request.
Local Proxy ARP	When this option is selected, local proxy ARP is enabled, and the interface can respond to an ARP request for a host other than itself. Unlike proxy ARP, local proxy ARP allows the interface to respond to ARP requests for a host that is on the same subnet as the host that sent the ARP request. This feature is useful when a host is not permitted to reply to an ARP request from another host in the same subnet, for example when using the protected ports feature.
Destination Unreachables	When this option is selected, the interface is allowed to send <i>ICMP</i> Destination Unreachable message to a host if the intended destination cannot be reached for some reason. If this option is clear, the interface will not send ICMP Destination Unreachable messages to inform the host about the error in reaching the intended destination.
ICMP Redirects	When this option is selected, the interface is allowed to send ICMP Redirect messages. The device sends an ICMP Redirect message on an interface only if ICMP Redirects are enabled both globally and on the interface. An ICMP Redirect message notifies a host when a better route to a particular destination is available on the network segment.

Table 267: Interface Configuration Fields (continued)

Field	Description
Secondary IP Address	To add a secondary IP address on the interface, click the + (plus) symbol in the header row and enter the address in the appropriate field in the Secondary IP Address Configuration window. You can add one or more secondary IP addresses to an interface only if the interface already has a primary IP address. To remove a configured secondary IP address, click the - (minus) symbol associated with the entry to remove. To remove all configured secondary IP addresses, click the - (minus) symbol in the header row.
Secondary Subnet Mask	The subnet mask associated with the secondary IP address. You can configure this field in the Secondary IP Address Configuration window.

Routing IP Loopback Configuration

Use the **Routing IP Loopback Configuration** page to configure the IP routing settings for each loopback interface.

To access this page, click **Routing > IP > Loopback Configuration** in the navigation menu.

Table 268: IP Loopback Configuration Fields

Field	Description
Interface	The menu contains all loopback interfaces that can be configured for routing. To configure routing settings for an interface, select it from the menu and then configure the rest of the settings on the page.
IP Address	The IP address of the loopback interface.
Subnet Mask	The IP subnet mask for the interface (also known as the network mask or netmask).
Secondary IP Address	To add a secondary IP address on the interface, click the + (plus) symbol in the header row and enter the address in the appropriate field in the Secondary IP Address Configuration window. You can add one or more secondary IP addresses to an interface only if the interface already has a primary IP address. To remove a configured secondary IP address, click the - (minus) symbol associated with the entry to remove. To remove all configured secondary IP addresses, click the - (minus) symbol in the header row.
Secondary Subnet Mask	The subnet mask associated with the secondary IP address. You can configure this field in the Secondary IP Address Configuration window.

Click **Refresh** to update the information on the screen.

Routing IP Statistics

The statistics reported on the **Routing IP Statistics** page are as specified in RFC 1213.

To access this page, click **Routing > IP > Statistics** in the navigation menu.

Table 269: IP Statistics Fields

Field	Description
IpInReceives	The total number of input datagrams received from interfaces, including those received in error.
IpInHdrErrors	The number of input datagrams discarded due to errors in their IP headers, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, errors discovered in processing their IP options, etc.
IpInAddrErrors	The number of input datagrams discarded because the IP address in their IP header's destination field was not a valid address to be received at this entity. This count includes invalid addresses (for example, 0.0.0.0) and addresses of unsupported Classes (for example, Class E). For entities which are not IP Gateways and therefore do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address.
IpForwDatagrams	The number of input datagrams for which this entity was not their final IP destination, as a result of which an attempt was made to find a route to forward them to that final destination. In entities which do not act as IP Gateways, this counter includes only those packets which were Source-Routed via this entity, and the Source-Route option processing was successful.
IpInUnknownProtos	The number of locally-addressed datagrams received successfully but discarded because of an unknown or unsupported protocol.
IpInDiscards	The number of input IP datagrams for which no problems were encountered to prevent their continued processing, but which were discarded (for example, for lack of buffer space). Note that this counter does not include any datagrams discarded while awaiting re-assembly.
IpInDelivers	The total number of input datagrams successfully delivered to IP user-protocols (including <i>ICMP</i>).
IpOutRequests	The total number of IP datagrams which local IP user-protocols (including <i>ICMP</i>) supplied to IP in requests for transmission. Note that this counter does not include any datagrams counted in ipForwDatagrams.
IpOutDiscards	The number of output IP datagrams for which no problem was encountered to prevent their transmission to their destination, but which were discarded (for example, for lack of buffer space). Note that this counter would include datagrams counted in ipForwDatagrams if any such packets met this (discretionary) discard criterion.
IpOutNoRoutes	The number of IP datagrams discarded because no route could be found to transmit them to their destination. Note that this counter includes any packets counted in ipForwDatagrams which meet this 'no-route' criterion. Note that this includes any datagrams which a host cannot route because all of its default gateways are down.
IpReasmTimeout	The maximum number of seconds which received fragments are held while they are awaiting reassembly at this entity.
IpReasmReqds	The number of IP fragments received which needed to be reassembled at this entity.
IpReasmOKs	The number of IP datagrams successfully re-assembled.
IpReasmFails	The number of failures detected by the IP re-assembly algorithm (for whatever reason: timed out, errors, etc.). Note that this is not necessarily a count of discarded IP fragments since some algorithms can lose track of the number of fragments by combining them as they are received.

Table 269: IP Statistics Fields (continued)

Field	Description
IpFragOKs	The number of IP datagrams that have been successfully fragmented at this entity.
IpFragFails	The number of IP datagrams that have been discarded because they needed to be fragmented at this entity but could not be, for example, because their Don't Fragment flag was set.
IpFragCreates	The number of IP datagram fragments that have been generated as a result of fragmentation at this entity.
IpRoutingDiscards	The number of routing entries which were chosen to be discarded even though they are valid. One possible reason for discarding such an entry could be to free-up buffer space for other routing entries.
IcmpInMsgs	The total number of ICMP messages which the entity received. Note that this counter includes all those counted by icmpInErrors.
IcmpInErrors	The number of ICMP messages which the entity received but determined as having ICMP-specific errors (bad ICMP checksums, bad length, etc.).
IcmpInDestUnreachs	The number of ICMP Destination Unreachable messages received.
IcmpInTimeExcds	The number of ICMP Time Exceeded messages received.
IcmpInParmProbs	The number of ICMP Parameter Problem messages received.
IcmpInSrcQuenchs	The number of ICMP Source Quench messages received.
IcmpInRedirects	The number of ICMP Redirect messages received.
IcmpInEchos	The number of ICMP Echo (request) messages received.
IcmpInEchoReps	The number of ICMP Echo Reply messages received.
IcmpInTimestamps	The number of ICMP Timestamp (request) messages received.
IcmpInTimestampReps	The number of ICMP Timestamp Reply messages received.
IcmpInAddrMasks	The number of ICMP Address Mask Request messages received.
IcmpInAddrMaskReps	The number of ICMP Address Mask Reply messages received.
IcmpOutMsgs	The total number of ICMP messages which this entity attempted to send. Note that this counter includes all those counted by icmpOutErrors.
IcmpOutErrors	The number of ICMP messages which this entity did not send due to problems discovered within ICMP such as a lack of buffers. This value should not include errors discovered outside the ICMP layer such as the inability of IP to route the resultant datagram. In some implementations there may be no types of error which contribute to this counter's value.
IcmpOutDestUnreachs	The number of ICMP Destination Unreachable messages sent.
IcmpOutTimeExcds	The number of ICMP Time Exceeded messages sent.
IcmpOutParmProbs	The number of ICMP Parameter Problem messages sent.
IcmpOutSrcQuenchs	The number of ICMP Source Quench messages sent.
IcmpOutRedirects	The number of ICMP Redirect messages sent. For a host, this object is always zero, since hosts do not send redirects.
IcmpOutEchos	The number of ICMP Echo (request) messages sent.

Table 269: IP Statistics Fields (continued)

Field	Description
IcmpOutEchoReps	The number of ICMP Echo Reply messages sent.
IcmpOutTimestamps	The number of ICMP Timestamp (request) messages.
IcmpOutTimestampReps	The number of ICMP Timestamp Reply messages sent.
IcmpOutAddrMasks	The number of ICMP Address Mask Request messages sent.
IcmpOutAddrMaskReps	The number of ICMP Address Mask Reply messages sent.

Router

The **Routing > Router** menu contains links to web pages that configure and display route tables.

Route Table Summary

The route table manager collects routes from multiple sources: static routes, [RIP \(Routing Information Protocol\)](#) routes, [BGP](#) routes, or local routes. The route table manager may learn multiple routes to the same destination from multiple sources. The **Route Table Summary** page lists all routes. The best routes table displays only the most preferred route to each destination.

To access this page, click **Routing > Router > Route Table** in the navigation menu.

Table 270: Route Table Fields

Field	Description
Network Address	The IP route prefix for the destination.
Subnet Mask	Also referred to as the subnet/network mask, this indicates the portion of the IP interface address that identifies the attached network.
Protocol	This field tells which protocol created the specified route. The possibilities are one of the following: <ul style="list-style-type: none"> • Local • Static • Default • RIP
Next Hop IP Address	The outgoing router IP address to use when forwarding traffic to the next router (if any) in the path towards the destination. The next router is always one of the adjacent neighbors or the IP address of the local interface for a directly-attached network.
Next Hop Interface	The outgoing interface to use when forwarding traffic to the destination. For a static reject route, the next hop is Null.
Best Route	Whether the route is the preferred route to the network. If the field is blank, a better route to the same network exists in the routing table.

Click **Refresh** to update the information on the screen.

Configured Route Summary

Use the **Configured Route Summary** page to create and display static routes.

To access this page, click **Routing > Router > Configured Routes** in the navigation menu.

Use the buttons to perform the following tasks:

- To configure a route, click **Add** and specify the desired settings in the available fields.
- To remove a configured route, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.

Table 271: Configured Routes Fields

Field	Description
Network Address	The IP route prefix for the destination.
Subnet Mask	Also referred to as the subnet/network mask, this indicates the portion of the IP interface address that identifies the attached network.
Next Hop IP	The next hop router address to use when forwarding traffic to the destination.
Next Hop Unit/Slot Port	The outgoing interface to use when forwarding traffic to the destination. For static reject routes it would be Null0.
Preference	The preferences configured for the added routes.

Adding and Removing Static Routes

The new route is added, and you are returned to the **Configured Routes** page.

- 1 Open the **Configured Routes** page.
- 2 Click **Add**.
The **Router Route Entry Configuration** window opens.
- 3 Next to **Route Type**, select one of the following options from the menu.
 - **Default**: Enter the default gateway address in the **Next Hop IP Address** field.
 - **Static**: Enter values for **Network Address**, **Subnet Mask**, **Next Hop IP Address**, and **Preference**.
 - **Static Reject**: Packets to these destinations will be dropped.



Note

The route type you select determines the fields available on the page. Some of the fields that [Table 272](#) describes are not available when configuring certain types of routes.

Table 272: Route Entry Create Fields

Field	Description
Network Address	Specify the IP route prefix for the destination from the drop-down menu. In order to create a route, a valid routing interface must exist and the next hop IP Address must be on the same network as the routing interface. Routing interfaces are created on the IP Interface Configuration page. Valid next hop IP Addresses can be viewed on the Route Table page.
Subnet Mask	Also referred to as the subnet/network mask, this indicates the portion of the IP interface address that identifies the attached network.

Table 272: Route Entry Create Fields (continued)

Field	Description
Protocol	This field tells which protocol created the specified route. Possible values are: <ul style="list-style-type: none"> • Local • Static • Default • RIP
Next Hop Slot/Port	The outgoing router interface to use when forwarding traffic to the destination.
Next Hop IP Address	The outgoing router IP address to use when forwarding traffic to the next router (if any) in the path towards the destination. The next router is always one of the adjacent neighbors or the IP address of the local interface for a directly attached network. When creating a route, the next hop IP must be on the same network as the routing interface. Valid next hop IP Addresses can be seen on the Route Table page.
Metric	Administrative cost of the path to the destination. If no value is entered, default is 1. The range is 0 — 255. This field is present only when creating a static route.
Preference	Specifies a preference value for the configured next hop.
Route Type	Specifies whether the route is to be a Default route or a Static route.

4 Click **Submit**.

To remove a configured route, click **Delete**.

IP Route Summary

The **IP Route Summary** page displays summary information about the entries in the IP routing table.

To access this page, click **Routing > Router > Summary** in the navigation menu.

Table 273: Summary Fields

Field	Description
Connected Routes	The total number of connected routes in the IP routing table.
Static Routes	The total number of static routes in the IP routing table.
RIP Routes	The total number of routes installed by the <i>RIP</i> protocol.
BGP Routes	The total number of routes installed by the <i>BGP</i> protocol.
External	The total number of external routes installed by the BGP protocol.
Internal	The total number of internal routes installed by the BGP protocol.
Local	The total number of local routes installed by the BGP protocol.
Reject Routes	The total number of reject routes installed by all protocols.
Total Routes	The total number of routes in the routing table.
Best Routes (High)	The number of best routes currently in the routing table. This number only counts the best route to each destination.

Table 273: Summary Fields (continued)

Field	Description
Alternate Routes	The number of alternate routes currently in the routing table. An alternate route is a route that was not selected as the best route to its destination.
Route Adds	The number of routes that have been added to the routing table.
Route Modifies	The number of routes that have been changed after they were initially added to the routing table.
Route Deletes	The number of routes that have been deleted from the routing table.
Unresolved Route Adds	The number of route adds that failed because none of the route's next hops were on a local subnet. Note that static routes can fail to be added to the routing table at startup because the routing interfaces are not yet up. This counter gets incremented in this case. The static routes are added to the routing table when the routing interfaces come up.
Invalid Route Adds	The number of routes that failed to be added to the routing table because the route was invalid. A log message is written for each of these failures.
Failed Route Adds	The number of routes that failed to be added to the routing table because of a resource limitation in the routing table.
Reserved Locals	The number of routing table entries reserved for a local subnet on a routing interface that is down. Space for local routes is always reserved so that local routes can be installed when a routing interface bounces.
Unique Next Hops (High)	The number of distinct next hops used among all routes currently in the routing table. These include local interfaces for local routes and neighbors for indirect routes.
Next Hop Groups (High)	The current number of next hop groups in use by one or more routes. Each next hop group includes one or more next hops.
ECMP Groups (High)	The number of next hop groups with multiple next hops.
ECMP Routes	The number of routes with multiple next hops currently in the routing table.
Truncated ECMP Routes	The number of <i>ECMP (Equal Cost Multi Paths)</i> routes that are currently installed in the forwarding table with just one next hop. The forwarding table may limit the number of ECMP routes or the number of ECMP groups. When an ECMP route cannot be installed because such a limit is reached, the route is installed with a single next hop.
ECMP Retries	The number of ECMP routes that have been installed in the forwarding table after initially being installed with a single next hop.
Routes with n Next Hops	The current number of routes with each number of next hops.
Clear Counters	This option resets to zero IPv4 routing table counters reported in this page. This option resets event counters only. Counters that report the current state of the routing table, such as the number of routes of each type, are not reset.

Click **Refresh** to update the information on the screen.

Configuring Routing Information Protocol (RIP)

RIP is part of the TCP/IP suite and maintains tables of all known destinations and the number of hops required to reach each. Using RIP, routers periodically exchange entire routing tables.

**Note**

This feature is available for 220 switches only.

Use the command-line interface to configure RIP. Refer to "Routing Information Protocol Commands" in *ExtremeSwitching 200 Series: Command Reference Guide*.

6 Managing Device Security

Port Access Control
RADIUS Settings
TACACS+ Settings
Authentication Manager

Use the features in the Security folder on the navigation menu to set management security parameters for port, user, and server security.

Port Access Control

In port-based authentication mode, when 802.1x is enabled globally and on the port, successful authentication of any one supplicant attached to the port results in all users being able to use the port without restrictions. At any given time, only one supplicant is allowed to attempt authentication on a port in this mode. Ports in this mode are under bidirectional control. This is the default authentication mode.

The 802.1X network has three components:

- **Authenticators** - Specifies the port that is authenticated before permitting system access.
- **Supplicants** - Specifies host connected to the authenticated port requesting access to the system services.
- **Authentication Server** - Specifies the external server, for example, the *RADIUS (Remote Authentication Dial In User Service)* server that performs the authentication on behalf of the authenticator, and indicates whether the user is authorized to access system services.

The Port Access Control folder contains links to the following pages that allow you to view and configure 802.1X features on the system.

Global Port Access Control Configuration

Use the **Port Based Access Control Configuration** page to enable or disable port access control on the system.

To access this page, click **Security > Port Access Control > Configuration** in the navigation menu.

Table 274: Port Access Control—Port Configuration Fields

Field	Description
Administrative Mode	Select Enable or Disable 802.1x mode on the switch. The default is Disable. This feature permits port-based authentication on the switch.
VLAN Assignment Mode	If enabled, when a supplicant is authenticated by a authentication server, the port that the supplicant is connected to is placed in a particular VLAN specified by the <i>RADIUS</i> server. VLAN Assignment mode controls if the switch is allowed to place a port in a RADIUS-assigned VLAN. A port's VLAN assignment is determined by the first supplicant that is authenticated on the port.
Dynamic VLAN Creation Mode	Select Enable to allow the switch to dynamically create a RADIUS-assigned VLAN if it does not already exist in the VLAN database.
Monitor Mode	The administrative mode of the Monitor Mode feature on the device. Monitor mode is a special mode that can be enabled in conjunction with port-based access control. Monitor mode provides a way for network administrators to identify possible issues with the port-based access control configuration on the device without affecting the network access to the users of the device. It allows network access even in cases where there is a failure to authenticate, but it logs the results of the authentication process for diagnostic purposes. If the device fails to authenticate a client for any reason (for example, RADIUS access reject from the RADIUS server, RADIUS timeout, or the client itself is 802.1X unaware), the client is authenticated and is undisturbed by the failure condition(s). The reasons for failure are logged and buffered into the local logging database for tracking purposes.
EAPOL Flood Mode	The administrative mode of the Extensible Authentication Protocol (EAP) over LAN (EAPOL) flood support on the device. EAPOL Flood Mode can be enabled when Admin Mode and Monitor Mode are disabled.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Port Access Control Port Summary

Use the **Port Access Control Port Summary** page to view summary information about the port-based authentication settings for each port.

To access this page, click **Security > Port Access Control > Port Summary** in the navigation menu.

Use the buttons to perform the following tasks:

- To change the port-based access control settings for a port, select the port to configure and click **Edit**. You are automatically redirected to the Port Access Control Port Configuration page for the selected port.
- To view additional information about the port-based access control settings for a port, select the port with the information to view and click **Details**. You are automatically redirected to the Port Access Control Port Details page for the selected port.

Table 275: Port Access Control—Port Summary Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.
PAE Capabilities	The Port Access Entity (PAE) role, which is one of the following: <ul style="list-style-type: none"> • Authenticator: The port enforces authentication and passes authentication information from a remote supplicant (similar to a client or host) to the authentication server. If the server successfully authenticates the supplicant, the port allows access. • Supplicant: The port must be granted permission by the authentication server before it can access the remote authenticator port.
Control Mode	The port-based access control mode configured on the port, which is one of the following: <ul style="list-style-type: none"> • Auto: The port is unauthorized until a successful authentication exchange has taken place. • Force Unauthorized: The port ignores supplicant authentication attempts and does not provide authentication services to the client. • Force Authorized: The port sends and receives normal traffic without client port-based authentication. • MAC-Based: This mode allows multiple supplicants connected to the same port to each authenticate individually. Each host connected to the port must authenticate separately in order to gain access to the network. The hosts are distinguished by their MAC addresses.
Operating Control Mode	The control mode under which the port is actually operating, which is one of the following: <ul style="list-style-type: none"> • Auto • Force Unauthorized • Force Authorized • MAC-Based • N/A <p>If the mode is N/A, port-based access control is not applicable to the port. If the port is in detached state it cannot participate in port access control. Additionally, if port-based access control is globally disabled, the status for all ports is N/A.</p>
PAE State	The current state of the authenticator PAE state machine, which is the 802.1X process that controls access to the port. The state can be one of the following: <ul style="list-style-type: none"> • Initialize • Disconnected • Connecting • Authenticating • Authenticated • Aborting • Held • ForceAuthorized • ForceUnauthorized

Table 275: Port Access Control—Port Summary Fields (continued)

Field	Description
Backend State	The current state of the back-end authentication state machine, which is the 802.1X process that controls the interaction between the 802.1X client on the local system and the remote authentication server. The state can be one of the following: <ul style="list-style-type: none"> • Request • Response • Success • Fail • Timeout • Initialize • Idle
Initialize (Icon)	Click the Initialize icon to reset the 802.1X state machine on the associated interface to the initialization state. Traffic sent to and from the port is blocked during the authentication process. This icon can be clicked only when the port is an authenticator and the operating control mode is Auto.
Re-Authenticate (Icon)	Click the Re-Authenticate icon to force the associated interface to restart the authentication process.

If you change any of the parameters, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must save the configuration.

Port Access Control Port Configuration

Use the **Port Access Control Port Configuration** page to enable and configure port access control on one or more ports.

To access this page, click **Security > Port Access Control > Port Configuration** in the navigation menu.

Use the buttons to perform the following tasks:

- To configure the port-based access control settings for one or more ports, select each port to configure and click **Edit**. The same settings are applied to all selected ports.
- To view additional information about the port-based access control settings for a port, select the port with the information to view and click **Details**.

Table 276: Port Access Control Port Configuration Fields

Field	Description
Interface	The interface with the settings to view or configure. If you have been redirected to this page, this field is read-only and displays the interface that was selected on the Port Access Control Port Summary page.
PAE Capabilities	<p>The Port Access Entity (PAE) role, which is one of the following:</p> <ul style="list-style-type: none"> • Authenticator: The port enforces authentication and passes authentication information from a remote supplicant (client or host) to the authentication server. If the server successfully authenticates the supplicant, the port allows access. • Supplicant: The port is connected to an authenticator port and must be granted permission by the authentication server before it can send and receive traffic through the remote port. <p>To change the PAE capabilities of a port, click the Edit icon associated with the field and select the desired setting from the menu in the Set PAE Capabilities window.</p>
Authenticator Options	The fields in this section can be changed only when the selected port is configured as an authenticator port (that is, the PAE Capabilities field is set to Authenticator).
Control Mode	<p>The port-based access control mode on the port, which is one of the following:</p> <ul style="list-style-type: none"> • Auto: The port is unauthorized until a successful authentication exchange has taken place. • Force Unauthorized: The port ignores supplicant authentication attempts and does not provide authentication services to the client. • Force Authorized: The port sends and receives normal traffic without client port-based authentication. • MAC-Based: This mode allows multiple supplicants connected to the same port to each authenticate individually. Each host connected to the port must authenticate separately in order to gain access to the network. The hosts are distinguished by their MAC addresses.
Quiet Period	The number of seconds that the port remains in the quiet state following a failed authentication exchange.
Transmit Period	The value, in seconds, of the timer used by the authenticator state machine on the port to determine when to send an EAPOL EAP Request/Identity frame to the supplicant.
Guest VLAN ID	The value, in seconds, of the timer used for guest VLAN authentication.
Unauthenticated VLAN ID	The VLAN ID of the unauthenticated VLAN. Hosts that fail the authentication might be denied access to the network or placed on a VLAN created for unauthenticated clients. This VLAN might be configured with limited network access. To set the unauthenticated VLAN ID, click the Edit icon associated with the field and specify the ID value in the available field. To reset the unauthenticated VLAN ID to the default value, click the Reset icon associated with the field and confirm the action.
Supplicant Timeout	The amount of time that the port waits for a response before retransmitting an EAP request frame to the client.
Server Timeout	The amount of time the port waits for a response from the authentication server.

Table 276: Port Access Control Port Configuration Fields (continued)

Field	Description
Maximum Requests	The maximum number of times that the port sends an EAP request frame (assuming that no response is received) to the client before restarting the authentication process.
Re-Authentication Period	The amount of time that clients can be connected to the port without being reauthenticated. If this field is disabled, connected clients are not forced to reauthenticate periodically. To change the value, click the Edit icon associated with the field and specify a value in the available field. To reset the reauthentication period to the default value, click the Reset icon associated with the field and confirm the action.
Maximum Users	The maximum number of clients supported on the port if the Control Mode on the port is MAC-based 802.1X authentication.
Supplicant Options	The fields in this section can be changed only when the selected port is configured as a supplicant port (that is, the PAE Capabilities field is set to Supplicant).
Control Mode	The port-based access control mode on the port, which is one of the following: <ul style="list-style-type: none"> • Auto: The port is in an unauthorized state until a successful authentication exchange has taken place between the supplicant port, the authenticator port, and the authentication server. • Force Unauthorized: The port is placed into an unauthorized state and is automatically denied system access. • Force Authorized: The port is placed into an authorized state and does not require client port-based authentication to be able to send and receive traffic.
User Name	The name the port uses to identify itself as a supplicant to the authenticator port. The menu includes the users that are configured for system management. When authenticating, the supplicant provides the password associated with the selected User Name.
Authentication Period	The amount of time the supplicant port waits to receive a challenge from the authentication server. If the configured Authentication Period expires, the supplicant retransmits the authentication request until it is authenticated or has sent the number of messages configured in the Maximum Start Messages field.
Start Period	The amount of time the supplicant port waits for a response from the authenticator port after sending a Start packet. If no response is received, the supplicant retransmits the Start packet.
Held Period	The amount of time the supplicant port waits before contacting the authenticator port after an active 802.1X session fails.
Maximum Start Messages	The maximum number of Start packets the supplicant port sends to the authenticator port without receiving a response before it considers the authenticator to be 802.1X-unaware.

Click **Refresh** to update the information on the screen.

Port Details

Use the **Port Access Control Port Details** page to view 802.1X information for a specific port.

To access this page, click **Security > Port Access Control > Port Details** in the navigation menu.

Table 277: Port Access Control Port Details Fields

Field	Description
Interface	The interface associated with the rest of the data on the page.
PAE Capabilities	The Port Access Entity (PAE) role, which is one of the following: <ul style="list-style-type: none"> • Authenticator: The port enforces authentication and passes authentication information from a remote supplicant (client or host) to the authentication server. If the server successfully authenticates the supplicant, the port allows access. • Supplicant: The port is connected to an authenticator port and must be granted permission by the authentication server before it can send and receive traffic through the remote port.
Authenticator Options	The fields in this section provide information about the settings that apply to the port when it is configured as an 802.1X authenticator.
Control Mode	The port-based access control mode on the port, which is one of the following: <ul style="list-style-type: none"> • Auto: The port is unauthorized until a successful authentication exchange has taken place. • Force Unauthorized: The port ignores supplicant authentication attempts and does not provide authentication services to the client. • Force Authorized: The port sends and receives normal traffic without client port-based authentication. • MAC-Based: This mode allows multiple supplicants connected to the same port to each authenticate individually. Each host connected to the port must authenticate separately in order to gain access to the network. The hosts are distinguished by their MAC addresses.
Quiet Period	The number of seconds that the port remains in the quiet state following a failed authentication exchange.
Transmit Period	The value, in seconds, of the timer used by the authenticator state machine on the port to determine when to send an EAPOL EAP Request/Identity frame to the supplicant.
Guest VLAN ID	The VLAN ID for the guest VLAN. The guest VLAN allows the port to provide a distinguished service to unauthenticated users. This feature provides a mechanism to allow users access to hosts on the guest VLAN.
Guest VLAN Period	The value, in seconds, of the timer used for guest VLAN authentication.
Unauthenticated VLAN ID	The VLAN ID of the unauthenticated VLAN. Hosts that fail the authentication might be denied access to the network or placed on a VLAN created for unauthenticated clients. This VLAN might be configured with limited network access.
Supplicant Timeout	The amount of time that the port waits for a response before retransmitting an EAP request frame to the client.
Server Timeout	The amount of time the port waits for a response from the authentication server.
Maximum Requests	The maximum number of times that the port sends an EAP request frame (assuming that no response is received) to the client before restarting the authentication process.

Table 277: Port Access Control Port Details Fields (continued)

Field	Description
Re-Authentication Period	The amount of time that clients can be connected to the port without being reauthenticated. If this field is disabled, connected clients are not forced to reauthenticate periodically.
Maximum Users	The maximum number of clients supported on the port if the Control Mode on the port is MAC-based 802.1X authentication.
Logical Port	The logical port number associated with the supplicant that is connected to the port.
Supplicant MAC Address	The MAC address of the supplicant that is connected to the port.
Authenticator PAE State	The current state of the authenticator PAE state machine, which is the 802.1X process that controls access to the port. The state can be one of the following: <ul style="list-style-type: none"> • Initialize • Disconnected • Connecting • Authenticating • Authenticated • Aborting • Held • ForceAuthorized • ForceUnauthorized
Backend Authentication State	The current state of the backend authentication state machine, which is the 802.1X process that controls the interaction between the 802.1X client on the local system and the remote authentication server. The state can be one of the following: <ul style="list-style-type: none"> • Request • Response • Success • Fail • Timeout • Initialize • Idle
VLAN Assigned	The ID of the VLAN the supplicant was placed in as a result of the authentication process.
VLAN Assigned Reason	The reason why the authenticator placed the supplicant in the VLAN. Possible values are: <ul style="list-style-type: none"> • RADIUS • Default • Not Assigned
Supplicant Options	The fields in this section provide information about the settings that apply to the port when it is configured as an 802.1X supplicant.

Table 277: Port Access Control Port Details Fields (continued)

Field	Description
<ul style="list-style-type: none"> Control Mode 	<p>The port-based access control mode on the port, which is one of the following:</p> <ul style="list-style-type: none"> Auto: The port is in an unauthorized state until a successful authentication exchange has taken place between the supplicant port, the authenticator port, and the authentication server. Force Unauthorized: The port is placed into an unauthorized state and is automatically denied system access. Force Authorized: The port is placed into an authorized state and does not require client port-based authentication to be able to send and receive traffic.
User Name	The name the port uses to identify itself as a supplicant to the authenticator port. The menu includes the users that are configured for system management. When authenticating, the supplicant provides the password associated with the selected User Name.
Authentication Period	The amount of time the supplicant port waits to receive a challenge from the authentication server. If the configured Authentication Period expires, the supplicant retransmits the authentication request until it is authenticated or has sent the number of messages configured in the Maximum Start Messages field.
Start Period	The amount of time the supplicant port waits for a response from the authenticator port after sending a Start packet. If no response is received, the supplicant retransmits the Start packet.
Held Period	The amount of time the supplicant port waits before contacting the authenticator port after an active 802.1X session fails.
Maximum Start Messages	The maximum number of Start packets the supplicant port sends to the authenticator port without receiving a response before it considers the authenticator to be 802.1X-unaware.

Click **Refresh** to update the information on the screen.

Statistics

Use the **Port Access Control Statistics** page to view information about the Extensible Authentication Protocol over LAN (EAPOL) frames and EAP messages sent and received by the local interfaces. To view additional per-interface EAPOL and EAP message statistics, select the interface with the information to view and click **Details**.

To access this page, click **Security > Port Access Control > Statistics** in the navigation menu.

Table 278: Port Access Control Statistics Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When viewing detailed information for an interface, this field identifies the interface being viewed.
PAE Capabilities	The Port Access Entity (PAE) role, which is one of the following: <ul style="list-style-type: none"> • Authenticator: The port enforces authentication and passes authentication information from a remote supplicant (similar to a client or host) to the authentication server. If the server successfully authenticates the supplicant, the port allows access. • Supplicant: The port must be granted permission by the authentication server before it can access the remote authenticator port.
EAPOL Frames Received	The total number of valid EAPOL frames received on the interface.
Last EAPOL Frame Version	The total number of EAPOL frames sent by the interface.
Last EAPOL Frame Source	After you click Details , a window opens and displays additional information about the EAPOL and EAP messages the interface sends and receives. The following information describes the additional fields that appear in the Details window. The fields this window displays depend on whether the interface is configured as an authenticator or supplicant, as noted in the applicable field descriptions.
EAPOL Start Frames Received	The total number of EAPOL-Start frames received on the interface. EAPOL-Start frames are sent by a supplicant to initiate the 802.1X authentication process when it connects to the interface. This field is displayed only if the interface is configured as an authenticator.
EAPOL Logoff Frames Received	The total number of EAPOL-Logoff frames received on the interface. EAPOL-Logoff frames are sent by a supplicant to indicate that it is disconnecting from the network, and the interface can return to the unauthorized state. This field is displayed only if the interface is configured as an authenticator.
EAP Response/ID Frames Received	The total number of EAP-Response Identity frames the interface has received. EAP-Response Identity frames are sent by a supplicant to provide user information that is used to for authentication. This field is displayed only if the interface is configured as an authenticator.
EAP Response Frames Received	The total number of EAP-Response frames the interface has received. EAP-Response frames are sent from a supplicant to an authentication server during the authentication process. This field is displayed only if the interface is configured as an authenticator.
EAP Request/ID Frames Transmitted	The total number of EAP-Request Identity frames the interface has sent. EAP-Request Identity frames are sent from an authenticator to a supplicant to request user information that is used to for authentication. This field is displayed only if the interface is configured as an authenticator.
EAPOL Start Frames Transmitted	The total number of EAPOL-Start frames the interface has sent to a remote authenticator. EAPOL-Start frames are sent by a supplicant to initiate the 802.1X authentication process when it connects to the interface. This field is displayed only if the interface is configured as a supplicant.
EAPOL Logoff Frames Transmitted	The total number of EAPOL-Logoff frames the interface has sent to a remote authenticator. EAPOL-Logoff frames are sent by a supplicant to indicate that it is disconnecting from the network, and the interface can return to the unauthorized state. This field is displayed only if the interface is configured as a supplicant.

Table 278: Port Access Control Statistics Fields (continued)

Field	Description
EAP Response/ID Frames Transmitted	The total number of EAP-Response Identity frames the interface has sent. EAP-Response Identity frames are sent by a supplicant to provide user information that is used to for authentication. This field is displayed only if the interface is configured as a supplicant.
EAP Request/ID Frames Received	The total number of EAP-Request Identity frames the interface has received. EAP-Request Identity frames are sent from an authenticator to a supplicant to request user information that is used to for authentication. This field is displayed only if the interface is configured as a supplicant.
EAP Request Frames Received	The total number of EAP-Request frames the interface has received. EAP-Request frames are sent from the authentication server to the supplicant during the authentication process. This field is displayed only if the interface is configured as a supplicant.
Invalid EAPOL Frames Received	The number of unrecognized EAPOL frames received on the interface.
EAPOL Length Error Frames Received	The number of EAPOL frames with an invalid packet body length received on the interface.
Clear (Button)	Resets all statistics counters to 0 for the selected interface or interfaces.

Click **Refresh** to update the information on the screen.

Client Summary

The **Port Access Control Client Summary** displays information about supplicant devices that are connected to the local authenticator ports. If there are no active 802.1X sessions, the table is empty. To view additional information about a supplicant, select the interface it is connected to and click **Details**.

To access this page, click **Security > Port Access Control > Client Summary** in the navigation menu.

Table 279: Port Access Control Client Summary Fields

Field	Description
Interface	The local interface associated with the rest of the data in the row. When viewing detailed information for an interface, this field identifies the interface being viewed.
Logical Interface	The logical port number associated with the supplicant that is connected to the port.
User Name	The name the client uses to identify itself as a supplicant to the authentication server.
Supp MAC Address	The MAC address of the supplicant that is connected to the port.
Session Time	The amount of time that has passed since the connected supplicant was granted access to the network through the authenticator port.
Filter ID	The policy filter ID assigned by the authenticator to the supplicant device.
VLAN ID	The ID of the VLAN the supplicant was placed in as a result of the authentication process.

Table 279: Port Access Control Client Summary Fields (continued)

Field	Description
	After you click Details , a window opens and displays additional information about the client. The following information describes the additional fields that appear in the window.
Session Timeout	The reauthentication timeout period set by the <u>RADIUS</u> server to the supplicant device.
Session Termination Action	The termination action set by the RADIUS server that indicates the action that will take place once the supplicant reaches the session timeout value.

Click **Refresh** to update the information on the screen.

Privileges Summary

Use the **Port Access Control Privileges Summary** page to grant or deny port access to users configured on the system. To change the access control privileges for one or more ports, select each interface to configure and click **Edit**. The same settings are applied to all selected interfaces.

To access this page, click **Security > Port Access Control > Privileges Summary** in the navigation menu.

Table 280: Port Access Control Privileges Summary Fields

Field	Description
Interface	The local interface associated with the rest of the data in the row. When configuring access information for one or more interfaces, this field identifies each interface being configured.
Users	The users that are allowed access to the system through the associated port. When configuring user access for a port, the Available Users field lists the users configured on the system that are denied access to the port. The users in the Selected Users field are allowed access. To move a user from one field to the other, click the user to move (or [Ctrl] + click to select multiple users) and click the appropriate arrow.

Click **Refresh** to update the information on the screen.

History Log Summary

Use the **Port Access Control History Log Summary** page to grant or deny port access to users configured on the system. To change the access control privileges for one or more ports, select each interface to configure and click **Edit**. The same settings are applied to all selected interfaces.

To access this page, click **Security > Port Access Control > History Log Summary** in the navigation menu.

Table 281: Port Access Control History Log Summary Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. Only interfaces that have entries in the log history are listed.
Time Stamp	The absolute time when the authentication event took place.
VLAN Assigned	The ID of the VLAN the supplicant was placed in as a result of the authentication process.
VLAN Assigned Reason	The reason why the authenticator placed the supplicant in the VLAN. Possible values are: <ul style="list-style-type: none"> • RADIUS • Unauth • Default • Not Assigned
Supp MAC Address	The MAC address of the supplicant that is connected to the port.
Filter Name	The policy filter ID assigned by the authenticator to the supplicant device.
Auth Status	The authentication status of the client or port.
Reason	The reason for the successful or unsuccessful authentication.

Click **Refresh** to update the information on the screen.

RADIUS Settings

[RADIUS](#) servers provide additional security for networks. The RADIUS server maintains a user database, which contains per-user authentication information. RADIUS servers provide a centralized authentication method for:

- Telnet Access
- Web Access
- Console to Switch Access
- Port Access Control (802.1X)

The RADIUS folder contains links to pages that help you view and configure system RADIUS settings.

RADIUS Configuration

Use the **RADIUS Configuration** page to view and configure various settings for the [RADIUS](#) servers configured on the system.

To access this page, click **Security > RADIUS > Configuration** in the navigation menu.

Table 282: RADIUS Configuration Fields

Field	Description
Max Number of Retransmits	The maximum number of times the RADIUS client on the device will retransmit a request packet to a configured RADIUS server after a response is not received. If multiple RADIUS servers are configured, the max retransmit value will be exhausted on the first server before the next server is attempted. A retransmit will not occur until the configured timeout value on that server has passed without a response from the RADIUS server. Therefore, the maximum delay in receiving a response from the RADIUS server equals the sum of (retransmit × timeout) for all configured servers. If the RADIUS request was generated by a user login attempt, all user interfaces will be blocked until the RADIUS application returns a response.
Timeout Duration	The number of seconds the RADIUS client waits for a response from the RADIUS server. Consideration to maximum delay time should be given when configuring RADIUS timeout and RADIUS max retransmit values.
Accounting Mode	Specifies whether the RADIUS accounting mode on the device is enabled or disabled.
NAS-IP Address	The network access server (NAS) IP address for the RADIUS server. To specify an address, click the Edit icon and enter the IP address of the NAS in the available field. The address should be unique to the NAS within the scope of the RADIUS server. The NAS IP address is used only in Access-Request packets. To reset the NAS IP address to the default value, click the Reset icon and confirm the action.

Use the buttons at the bottom of the page to perform the following actions:

- Click **Refresh** to update the page with the most current information.
- If you make changes to the page, click **Submit** to apply the changes to the system.

Named Server Status

The **RADIUS Named Server Status** page shows summary information about the *RADIUS* servers configured on the system.

To access this page, click **Security > RADIUS > Named Server** in the navigation menu.

Use the buttons to perform the following tasks:

- To add a RADIUS authentication server to the list of servers the RADIUS client can contact, click **Add**.
- To change the settings for a configured RADIUS server, select the entry to modify and click **Edit**. You cannot change the IP address or host name for a server after it has been added.
- To remove a configured RADIUS server from the list, select the entry to delete and click **Remove**. You must confirm the action before the entry is deleted.

Table 283: RADIUS Server Status Fields

Field	Description
Current	An asterisk (*) in the column Indicates that the server is the current server for the authentication server group. If no asterisk is present, the server is a backup server. If more than one RADIUS server is configured with the same name, the switch selects one of the servers to be the current server from the group of servers with the same name. When the switch sends a RADIUS request to the named server, the request is directed to the server selected as the current server. Initially the primary server is selected as the current server. If the primary server fails, one of the other servers becomes the current server.
RADIUS Server Host Address	Shows the IP address of the RADIUS server.
RADIUS Server Name	Shows the RADIUS server name. Multiple RADIUS servers can have the same name. In this case, RADIUS clients can use RADIUS servers with the same name as backups for each other.
Port Number	Identifies the authentication port the server uses to verify the RADIUS server authentication. The port is a UDP port.
Server Type	Shows whether the server is a Primary or Secondary server.
Secret Configured	Whether the shared secret for this server has been configured.
Message Authenticator	Shows whether the message authenticator attribute for the selected server is enabled or disabled.

Click **Refresh** to update the page with the most current information.

Server Statistics

Use the **RADIUS Server Statistics** page to view statistical information for each RADIUS server configured on the system.

To access this page, click **Security > RADIUS > Statistics** in the navigation menu.

Table 284: RADIUS Server Statistics Fields

Field	Description
IP Address/Host Name	The IP address or host name of the RADIUS server associated with the rest of the data in the row. When viewing the detailed statistics for a RADIUS server, this field identifies the RADIUS server.
Round Trip Time	The time interval, in hundredths of a second, between the most recent Access-Reply/Access-Challenge and the Access-Request that matched it from the RADIUS authentication server.
Access Requests	The number of RADIUS Access-Request packets sent to the server. This number does not include retransmissions.
Access Rejects	The number of RADIUS Access-Reject packets, including both valid and invalid packets, that were received from the server.
Pending Requests	The number of RADIUS Access-Request packets destined for the server that have not yet timed out or received a response.
Timeouts	The number of times a response was not received from the server within the configured timeout value.

Table 284: RADIUS Server Statistics Fields (continued)

Field	Description
Packets Dropped	The number of RADIUS packets received from the server on the authentication port and dropped for some other reason.
Access Retransmissions	The number of RADIUS Access-Request packets that had to be retransmitted to the server because the initial Access-Request packet failed to be successfully delivered.
Access Accepts	The number of RADIUS Access-Accept packets, including both valid and invalid packets, that were received from the server.
Access Challenges	The number of RADIUS Access-Challenge packets, including both valid and invalid packets, that were received from the server.
Malformed Access Responses	The number of malformed RADIUS Access-Response packets received from the server. Malformed packets include packets with an invalid length. Bad authenticators, signature attributes, and unknown types are not included as malformed access responses.
Bad Authenticators	The number of RADIUS Access-Response packets containing invalid authenticators or signature attributes received from the server.
Unknown Types	The number of RADIUS packets of unknown type which were received from the server on the authentication port.

Click **Refresh** to update the page with the most current information.

RADIUS Accounting Server Status

The **RADIUS Accounting Server Status** page shows summary information about the [RADIUS](#) accounting servers configured on the system.

To access this page, click **Security > RADIUS > Accounting Server** in the navigation menu.

Use the buttons to perform the following tasks:

- To add a RADIUS accounting server to the list of servers the RADIUS client can contact, click **Add**.
- To change the settings for a configured RADIUS accounting server, select the entry to modify and click **Edit**. You cannot change the IP address or host name for a server after it has been added.
- To remove a configured RADIUS accounting server from the list, select the entry to delete and click **Remove**. You must confirm the action before the entry is deleted.

Table 285: RADIUS Accounting Server Status Fields

Field	Description
IP Address/Host Name	The IP address or host name of the RADIUS accounting server. Host names must be resolvable by DNS and are composed of a series of labels separated by dots.
Server Name	The name of the RADIUS accounting server. RADIUS servers that are configured with the same name are members of the same named RADIUS server group. RADIUS accounting servers in the same group serve as backups for each other.
Port Number	The UDP port on the RADIUS accounting server to which the local RADIUS client sends request packets.

Table 285: RADIUS Accounting Server Status Fields (continued)

Field	Description
Secret Configured	Whether the shared secret for this server has been configured.
Secret	The shared secret text string used for authenticating and encrypting all RADIUS communications between the RADIUS client on the device and the RADIUS accounting server. The secret specified in this field must match the shared secret configured on the RADIUS accounting server.

Click **Refresh** to update the page with the most current information.

Accounting Statistics

Use the **RADIUS Accounting Statistics** page to view statistical information for each *RADIUS* server configured on the system.

To access this page, click **Security > RADIUS > Accounting Statistics** in the navigation menu.

Table 286: RADIUS Accounting Statistics Fields

Field	Description
IP Address/Host Name	The IP address or host name of the RADIUS accounting server associated with the rest of the data in the row. When viewing the detailed statistics for a RADIUS accounting server, this field identifies the server.
Round Trip Time	Displays the time interval, in hundredths of a second, between the most recent Accounting-Response and the Accounting-Request that matched it from this RADIUS accounting server.
Accounting Requests	The number of RADIUS Accounting-Request packets sent to this server. This number does not include retransmissions.
Pending Requests	The number of RADIUS Accounting-Request packets destined for the server that have not yet timed out or received a response.
Timeouts	The number of times a response was not received from the server within the configured timeout value.
Packets Dropped	The number of RADIUS packets received from the server on the accounting port and dropped for some other reason.
Accounting Retransmissions	The number of RADIUS Accounting-Request packets retransmitted to the server.
Accounting Responses	The number of RADIUS packets received on the accounting port from the server.
Timeouts	The number of accounting timeouts to this server.
Malformed Access Responses	The number of malformed RADIUS Accounting-Response packets received from the server. Malformed packets include packets with an invalid length. Bad authenticators and unknown types are not included as malformed accounting responses.
Bad Authenticators	The number of RADIUS Accounting-Response packets that contained invalid authenticators received from the accounting server.
Unknown Types	The number of RADIUS packets of unknown type which were received from the server on the accounting port.

RADIUS Clear Statistics

Use the **RADIUS Clear Statistics** page to reset all *RADIUS* authentication and accounting statistics to zero.

To access this page, click **Security > RADIUS > Clear Statistics** in the navigation menu.

To clear all statistics for the RADIUS authentication and accounting server, click **Reset**. After you confirm the action, the statistics on both the **RADIUS Server Statistics** and **RADIUS Accounting Server Statistics** pages are reset.

RADIUS Source Interface Configuration

Use the **RADIUS Source Interface Configuration** page to specify the physical or logical interface to use as the *RADIUS* client source interface. When an IP address is configured on the source interface, this address is used for all RADIUS communications between the local RADIUS client and the remote RADIUS server. The IP address of the designated source interface is used in the IP header of RADIUS management protocol packets. This allows security devices, such as firewalls, to identify all source packets coming from a specific device.

To access this page, click **Security > RADIUS > Source Interface Configuration** in the navigation menu.

Table 287: RADIUS Accounting Statistics Fields

Field	Description
Type	The type of interface to use as the source interface: <ul style="list-style-type: none"> • None: The primary IP address of the originating (outbound) interface is used as the source address. • Interface: The primary IP address of a physical port is used as the source address. • Loopback: The primary IP address of the loopback interface is used as the source address. A loopback is always reachable, as long as any routing interface is up. • VLAN: The primary IP address of a VLAN routing interface is used as the source address.
Interface	When the selected Type is Interface, select the physical port to use as the source interface.
Loopback Interface	When the selected Type is Loopback, select the loopback interface to use as the source interface.
VLAN ID	When the selected Type is VLAN, select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.

Click **Refresh** to update the page with the most current information.

TACACS+ Settings

To access the **TACACS+ Configuration** page, click **Security > TACACS+ > Configuration** in the navigation menu.

Table 288: TACACS+ Configuration Fields

Field	Description
Key String	Specifies the authentication and encryption key for TACACS+ communications between the device and the TACACS+ server. The key must match the key configured on the TACACS+ server.
Connection Timeout	The maximum number of seconds allowed to establish a TCP connection between the device and the TACACS+ server.

Click **Refresh** to update the page with the most current information.

If you make any changes to the page, click **Submit** to apply the changes to the system.

TACACS+ Server Summary

Use the **TACACS+ Server Summary** page to view and configure information about the TACACS+ server(s).

To access this page, click **Security > TACACS+ > Server Summary** in the navigation menu.

Use the buttons to perform the following tasks:

- To add a TACACS+ Server to the list of servers the TACACS+ client can contact, click **Add**. If maximum number of server is added, the button will be disabled
- To edit a configured TACACS+ server from the list, select the entry and click **Edit**.
- To remove a configured TACACS+ server from the list, select the entry to delete and click **Remove**. You must confirm the action before the entry is deleted.

Table 289: TACACS+ Server Summary Fields

Field	Description
Server	Specifies the TACACS+ Server IP address or Hostname.
Priority	Specifies the order in which the TACACS+ servers are used.
Port	Specifies the authentication port.
Key String	Specifies the authentication and encryption key for TACACS+ communications between the device and the TACACS+ server. The key must match the encryption used on the TACACS+ server.
Connection Timeout	The amount of time that passes before the connection between the device and the TACACS+ server time out.

Click **Refresh** to update the page with the most current information.

TACACS+ Server Configuration

Use the **TACACS+ Server Configuration** page to view and configure information about the TACACS+ server(s).

To access this page, click **Security > TACACS+ > Server Configuration** in the navigation menu.

Table 290: TACACS+ Server Configuration Fields

Field	Description
Server	The TACACS+ Server IP address or Hostname.
Priority	The order in which the TACACS+ servers are used.
Port	The authentication port.
Key String	The authentication and encryption key for TACACS+ communications between the device and the TACACS+ server. The key must match the encryption used on the TACACS+ server.
Connection Timeout	The amount of time that passes before the connection between the device and the TACACS+ server time out.

Click **Refresh** to update the page with the most current information.

If you make any changes to the page, click **Submit** to apply the changes to the system.

TACACS+ Source Interface Configuration

Use the **TACACS+ Source Interface Configuration** page to specify the physical or logical interface to use as the **TACACS+** client source interface. When an IP address is configured on the source interface, this address is used for all TACACS+ communications between the local TACACS+ client and the remote TACACS+ server. The IP address of the designated source interface is used in the IP header of TACACS+ management protocol packets. This allows security devices, such as firewalls, to identify all source packets coming from a specific device.

To access this page, click **Security > TACACS+ > Source Interface Configuration** in the navigation menu.

Table 291: TACACS+ Source Interface Configuration Fields

Field	Description
Type	The type of interface to use as the source interface: <ul style="list-style-type: none"> • None: The primary IP address of the originating (outbound) interface is used as the source address. • Interface: The primary IP address of a physical port is used as the source address. • VLAN: The primary IP address of a VLAN routing interface is used as the source address. • Network: The network source IP is used as the source address. • Service Port: The management port source IP is used as the source address.
Interface	When the selected Type is Interface, select the physical port to use as the source interface.
VLAN ID	When the selected Type is VLAN, select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.

Click **Refresh** to update the page with the most current information.

If you make any changes to the page, click **Submit** to apply the changes to the system.

Authentication Manager

The Authentication Manager feature allows you to configure the authentication methods used on the individual interface.

Authentication Manager Configuration

Use the **Authentication Manager Configuration** page to control the administrative mode of the Authentication Manager feature, which enables configuration of the sequence and priority of the authentication methods per interface.

To access this page, click **Security > Authentication Manager > Configuration** in the navigation menu.

Table 292: RADIUS Configuration Fields

Field	Description
Admin Mode	The administrative mode of the Authentication Manager feature. When Authentication Manager is enabled globally, the authentication methods configured on an interface are executed in the order configured as the device attempts to authenticate clients on that interface.

Use the buttons at the bottom of the page to perform the following actions:

- Click **Submit** to apply the settings to the running configuration and cause the change to take effect. These changes will not be retained across a power cycle unless a **Save configuration** is performed.
- Click **Refresh** to display the latest information from the switch.
- Click **Cancel** to cancel the change.

Authentication Tiering

Use the **Authentication Tiering** page to configure the sequence and priority of the authentication methods for the interfaces on the device. When Authentication Manager is enabled globally, the authentication methods configured on an interface are executed in the order configured as the device attempts to authenticate clients on that interface. The default method order is Dot1x, MAC Authentication Bypass (MAB).

To access this page, click **Security > Authentication Manager > Authentication Tiering** in the navigation menu.

Table 293: Authentication Tiering Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When editing information for one or more interfaces, this field identifies the interfaces that are being configured.
Configured Order	The order in which the authentication methods are used to authenticate a client connected to an interface, which can be one or more of the following: <ul style="list-style-type: none"> • Dot1x: The port-based authentication method. • MAB: MAC Authentication Bypass method that uses the MAC address of the client to determine the kind of network access to provide.
Enabled Order	The methods from the list of authentication methods configured on an interface which are administratively enabled in the device.
Configured Priority	The priority of the authentication methods. The default priority of a method is equivalent to its position in the order of the authentication list configured per interface. If the priority of the methods is changed, all clients authenticated using a lower priority method are forced to re-authenticate.
Enabled Priority	The methods from the list of authentication method priorities configured on an interface which are administratively enabled in the device.
Authenticated Clients	Number of clients authenticated on an interface.
Re-Authentication Timer	Interval, in seconds, after which an attempt is made to authenticate an unauthorized port.

Use the buttons at the bottom of the page to perform the following actions:

- Click **Refresh** to display the latest information from the switch.
- Click **Edit** to configure the settings for one or more interfaces, select each entry to modify. The settings are applied to all selected interfaces.

Authenticated Clients

Use the **Authentication Clients** page to view information about the clients connected on the interfaces. If there are no clients connected, the table is empty.

To access this page, click **Security > Authentication Manager > Authentication Clients** in the navigation menu.

Table 294: Authentication Clients Fields

Field	Description
Interface	The local interface associated with the rest of the data in the row.
Logical Interface	The logical port number associated with the client that is connected to the port.
Client MAC Address	The MAC address of the client that is connected to the port.

Table 294: Authentication Clients Fields (continued)

Field	Description
Authenticated Method	The authentication method used to authenticate a client connected to an interface, which can be one of the following: <ul style="list-style-type: none"> • Dot1x: The port-based authentication method. • MAB: MAC Authentication Bypass method that uses the MAC address of the client to determine the kind of network access to provide.
Authentication State	The current client authentication state, which can be one of the following: <ul style="list-style-type: none"> • Success: Indicates authentication succeeded. • Failure: Indicates authentication failed.
Authentication Status	The client authentication status, which can be one of the following: <ul style="list-style-type: none"> • Authorized: Indicates client is authorized on the port. • Unauthorized: Indicates client is not authorized on the port.

Click **Refresh** to display the latest information from the switch.

Authentication Statistics

Use the **Authentication Statistics** page to view information about the Authentication Manager client authentication attempts and failures per interface.

To access this page, click **Security > Authentication Manager > Statistics** in the navigation menu.

Table 295: Authentication Statistics Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.
Dot1x Attempts	The number of attempts made to authenticate a client using the Dot1x authentication method.
Dot1x Failures	The number of attempts that failed when Dot1x method is used for client authentication.
MAB Attempts	The number of attempts made to authenticate a client using the MAC Authentication Bypass (MAB) authentication method.
MAB Failures	The number of attempts that failed when MAB method is used for client authentication.
Captive Portal Attempts	The number of attempts made to authenticate a client using the Captive Portal authentication method. Note: Captive portal is not supported in this version of the product.
Captive Portal Failures	The number of attempts that failed when Captive Portal method is used for client authentication. Note: Captive portal is not supported in this version of the product.

Use the buttons at the bottom of the page to perform the following actions:

- Click **Refresh** to display the latest information from the switch.
- Click **Clear** to reset all statistics counters to 0 for the selected interfaces.

Authentication History

Use the **Authentication History** page to view the Authentication Manager history log per interface.

To access this page, click **Security > Authentication Manager > History** in the navigation menu.

Table 296: Authentication History Fields

Field	Description
Interface	The menu contains all interfaces in the device. To view the history log on a specific interface, select the interface from the menu.
Time Stamp	The absolute time when the authentication event took place.
MAC Address	The MAC address of the client that is connected to the port.
Authentication Status	The client authentication status, which can be one of the following: <ul style="list-style-type: none"> • Authorized: Indicates client is authorized on the port. • Unauthorized: Indicates client is not authorized on the port.
Authenticated Method	The authentication method used to authenticate a client connected to an interface, which can be one of the following: <ul style="list-style-type: none"> • Dot1x: The port-based authentication method. • MAB: MAC Authentication Bypass method that uses the MAC address of the client to determine the kind of network access to provide.

Use the buttons at the bottom of the page to perform the following actions:

- Click **Refresh** to display the latest information from the switch.
- Click **Clear** to clear the Authentication Manager history log on the selected interface.

7 Configuring IPv6

Global Configuration

IPv6 is the next generation of the Internet Protocol. With 128-bit addresses, versus 32-bit addresses for IPv4, IPv6 solves the address depletion issues seen with IPv4 and removes the requirement for Network Address Translation (NAT), which is used in IPv4 networks to reduce the number of globally unique IP addresses required for a given network. Its aggregate addresses can dramatically reduce the size of the global routing table through well known address combinations. Security is more integrated and network configuration is simplified yet more flexible.

IPv6 can coexist with IPv4. As with IPv4, IPv6 routing can be enabled on physical and VLAN interfaces. Each Layer 3 routing interface can be used for IPv4, IPv6, or both. IP protocols running over Layer 3 (for example, UDP and TCP) do not change with IPv6. For this reason, a single CPU stack is used for transport of both IPv4 and IPv6, and a single sockets interface provides access to both. Routing protocols are capable of computing routes for one or both IP versions.



Note

CLI commands are not available for all the IPv6 pages.

Global Configuration

Use the **IPv6 Network Connectivity** page to configure and view IPv6 information on the network interface. The network interface is the logical interface that allows remote management of the device via any of the front-panel switch ports. To enable management of the device over an IPv6 network by using a web browser, *SNMP (Simple Network Management Protocol)*, Telnet, or SSH, you must first configure the device with the appropriate IPv6 information. The configuration parameters associated with the network interface do not affect the configuration of the front-panel ports through which traffic is switched or routed.

To access this page, click **System > Connectivity > IPv6** in the navigation menu.

Table 297: IPv6 Network Connectivity Fields

Field	Description
IPv6 Mode	Enables or disables the IPv6 administrative mode on the network interface.
Network Configuration Protocol	Specify whether the device should attempt to acquire network information from a DHCPv6 server. Selecting None disables the DHCPv6 client on the network interface.

Table 297: IPv6 Network Connectivity Fields (continued)

Field	Description
IPv6 Stateless Address AutoConfig Mode	<p>Sets the IPv6 stateless address autoconfiguration mode on the network interface.</p> <ul style="list-style-type: none"> • Enabled: The network interface can acquire an IPv6 address through IPv6 Neighbor Discovery Protocol (NDP) and through the use of Router Advertisement messages. • Disabled: The network interface will not use the native IPv6 address autoconfiguration features to acquire an IPv6 address.
DHCPv6 Client DUID	The client identifier used by the DHCPv6 client (if enabled) when sending messages to the DHCPv6 server.
IPv6 Gateway	The default gateway for the IPv6 network interface. To configure this field, click the Edit icon in the row. To reset the field to the default value, click the Reset icon in the row.
Static IPv6 Addresses	<p>Lists the manually configured static IPv6 addresses on the network interface. Use the buttons available in this table to perform the following tasks:</p> <ul style="list-style-type: none"> • To add an entry to the list, click the + (plus) button to open the Add IPv6 Address dialog and provide the following: <ul style="list-style-type: none"> • New IPv6 Address: Specify the IPv6 address to add to the interface. • EUI Flag: Select this option to enable the Extended Universal Identifier (EUI) flag for IPv6 address, or clear the option to omit the flag. • To delete an entry from the list, click the – (minus) button associated with the entry to remove. • To delete all entries from the list, click the – (minus) button in the heading row.
Dynamic IPv6 Addresses	Lists the IPv6 addresses on the network interface that have been dynamically configured through IPv6 autoconfiguration or DHCPv6.
Default IPv6 Routers	Lists the IPv6 address of each default router that has been automatically configured through IPv6 router discovery.

Click **Refresh** to update the information on the screen.

8 Configuring Quality of Service

Configuring Access Control Lists
Configuring Auto VoIP
Configuring Class of Service
Configuring Diffserv

This section gives an overview of QoS (Quality of Service) and explains the QoS features available from the QoS navigation menu.

In a typical switch, each physical port consists of one or more queues for transmitting packets on the attached network. Multiple queues per port are often provided to give preference to certain packets over others based on user-defined criteria. When a packet is queued for transmission in a port, the rate at which it is serviced depends on how the queue is configured and possibly the amount of traffic present in the other queues of the port. If a delay is necessary, packets get held in the queue until the scheduler authorizes the queue for transmission. As queues become full, packets have no place to be held for transmission and get dropped by the switch.

QoS is a means of providing consistent, predictable data delivery by distinguishing between packets that have strict timing requirements from those that are more tolerant of delay. Packets with strict timing requirements are given “special treatment” in a QoS capable network. With this in mind, all elements of the network must be QoS-capable. The presence of at least one node which is not QoS-capable creates a deficiency in the network path and the performance of the entire packet flow is compromised.

Configuring Access Control Lists

An ACL (Access Control List) ensures that only authorized users have access to specific resources while blocking off any unwarranted attempts to reach network resources. ACLs are used to provide traffic flow control, restrict contents of routing updates, decide which types of traffic are forwarded or blocked, and above all provide security for the network. 200 Series software supports IPv4, IPv6, and MAC ACLs. The total number of MAC and IP ACLs supported by 200 Series software is platform-specific.

You first create an IPv4-based, IPv6-based, or MAC-based rule and assign a unique ACL ID. Then, you define the rules, which can identify protocols, source and destination IP and MAC addresses, and other packet-matching criteria. Finally, you use the ID number to assign the ACL to a port or to a VLAN interface.

IP Access Control Lists

An IP ACL allows network managers to define classification actions and rules for specific ports. ACLs are composed of ACE (Access Control Entry), or rules, that consist of the filters that determine traffic classifications. The total number of rules that can be defined for each ACL is platform-specific. These rules are matched sequentially against a packet. When a packet meets the match criteria of a rule, the

specified rule action (Permit/Deny) is taken, including dropping the packet or disabling the port, and the additional rules are not checked for a match. For example, a network administrator defines an ACL rule that says port number 20 can receive TCP packets. However, if a UDP packet is received, the packet is dropped.

The IP Access Control List folder contains links to web pages that allow you to configure and view IP ACLs.

To configure an IP ACL:

- 1 Use the [IP ACL Configuration](#) on page 299 page to define the IP ACL type and assign an ID to it.
- 2 Use the [Access Control List Interface Summary](#) on page 307 page to create rules for the ACL.
- 3 Use the [Access Control List Configuration](#) on page 300 page to view the configuration.

IP ACL Configuration

Use the **IP ACL Configuration** page to add or remove IP-based ACLs. On this menu the interfaces to which an IP ACL applies must be specified, as well as whether it applies to inbound or outbound traffic. Rules for the IP ACL are specified/created using the [Access Control List Interface Summary](#) on page 307.

To access this page page, click **QoS > Access Control Lists > Summary** in the navigation menu.

Use the buttons at the bottom of the page to perform the following tasks:

- To add an ACL, click **Add** and configure the ACL type and ID.
- To remove one or more configured ACLs, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.
- To configure rules for an ACL, select the ACL to configure and click **Edit**. You are redirected to the Access Control List Configuration page for the selected ACL.

Table 298: Access List Summary Fields

Field	Description
ACL Identifier	The name or number that identifies the ACL. The permitted identifier depends on the ACL type. Standard and Extended IPv4 ACLs use numbers within a set range, and Named IPv4 and MAC ACLs use alphanumeric characters.
ACL Type	The type of ACL. The ACL type determines the criteria that can be used to match packets. The type also determines which attributes can be applied to matching traffic. IPv4 ACLs classify Layer 3 and Layer 4 IPv4 traffic, IPv6 ACLs classify Layer 3 and Layer 4 IPv6 traffic, and MAC ACLs classify Layer 2 traffic. The ACL types are as follows: <ul style="list-style-type: none"> • IPv4 Standard: Match criteria is based on the source address of IPv4 packets. • IPv4 Extended: Match criteria can be based on the source and destination addresses, source and destination Layer 4 ports, and protocol type of IPv4 packets. • IPv4 Named: Match criteria is the same as IPv4 Extended ACLs, but the ACL ID can be an alphanumeric name instead of a number. • IPv6 Named: Match criteria can be based on information including the source and destination IPv6 addresses, source and destination Layer 4 ports, and protocol type within IPv6 packets. • Extended MAC: Match criteria can be based on the source and destination MAC addresses, 802.1p user priority, VLAN ID, and EtherType value within Ethernet frames.
Rules Used	The number of rules currently configured for the ACL
Direction	Whether the packet is checked against the rules in an ACL when it is received on an interface (Inbound) or after it has been received, routed, and is ready to exit an interface (Outbound).
Interface	The interface(s) to which the ACL has been applied.
VLAN	Each VLAN to which the ACL has been applied.

Access Control List Configuration

Use the **Access Control List Configuration** page to configure rules for each existing ACL on the system and to view summary information about the rules that have been added to an ACL. Each ACL rule is configured to match one or more aspects of traffic on the network. When a packet matches the conditions in a rule, it is handled according to the configured action (permit or deny) and attributes. Each ACL can have multiple rules, but the final rule for every ACL is an implicit deny all rule. For each rule, a packet must match all the specified criteria in order for the specified rule action (Permit/Deny) to take place.

To access this page, click **QoS > Access Control Lists > Configuration** in the navigation menu.

Use the buttons to perform the following tasks:

- To add an Access List Rule entry, select the ID of the ACL that will include the rule from the ACL Identifier menu. Then, click **Add Rule** and configure the rule criteria and attributes. New rules cannot be created if the maximum number of rules has been reached.
- To remove the most recently configured rule for an ACL, select the ID of the appropriate ACL from the ACL Identifier menu and click **Remove Last Rule**. You must confirm the action before the entry is deleted.

- To resequence rules for an ACL, select the ID of the appropriate ACL from the ACL Identifier menu and click **Resequence Rules**.

Table 299: IP ACL Configuration Fields

Field	Description
ACL Identifier	The menu contains the ID for each ACL that exists on the system. Before you add or remove a rule, you must select the ID of the ACL from the menu. For ACLs with alphanumeric names, click the Edit icon to change the ACL ID. The ID of a Named IPv4 ACL must begin with a letter, and not a number. The ACL identifier for IPv4 Standard and IPv4 Extended ACLs cannot be changed.
Sequence Number	The number that indicates the position of a rule within the ACL. If the sequence number is not specified during rule creation, the rule is automatically assigned a sequence number after it is successfully added to the ACL. The rules are displayed based on their position within the ACL, but can also be renumbered. Packets are checked against the rule criteria in order, from the lowest-numbered rule to the highest. When the packet matches the criteria in a rule, it is handled according to the rule action and attributes. If no rule matches a packet, the packet is discarded based on the implicit deny all rule, which is the final rule in every ACL.
ACL Type	The type of ACL. The ACL type determines the criteria that can be used to match packets. The type also determines which attributes can be applied to matching traffic. IPv4 ACLs classify Layer 3 and Layer 4 IPv4 traffic, IPv6 ACLs classify Layer 3 and Layer 4 IPv6 traffic, and MAC ACLs classify Layer 2 traffic. The ACL types are as follows: <ul style="list-style-type: none"> IPv4 Standard: Match criteria is based on the source address of IPv4 packets. IPv4 Extended: Match criteria can be based on the source and destination addresses, source and destination Layer 4 ports, and protocol type of IPv4 packets. IPv4 Named: Match criteria is the same as IPv4 Extended ACLs, but the ACL ID can be an alphanumeric name instead of a number. IPv6 Named: Match criteria can be based on information including the source and destination IPv6 addresses, source and destination Layer 4 ports, and protocol type within IPv6 packets. Extended MAC: Match criteria can be based on the source and destination MAC addresses, 802.1p user priority, VLAN ID, and EtherType value within Ethernet frames.
Status	Whether the ACL is active. If the ACL is a time-based ACL that includes a time range, the ACL is active only during the periods specified within the time range. If an ACL does not include a time range, the status is always active.
Action	The action to take when a packet or frame matches the criteria in the rule: <ul style="list-style-type: none"> Permit: The packet or frame is forwarded. Deny: The packet or frame is dropped. <p>When configuring ACL rules in the Add Access Control List Rule window, the selected action determines which fields can be configured. Not all fields are available for both Permit and Deny actions.</p>
Match Conditions	The criteria used to determine whether a packet or frame matches the ACL rule.

Table 299: IP ACL Configuration Fields (continued)

Field	Description
Rule Attributes	Each action — beyond the basic Permit and Deny actions — to perform on the traffic that matches the rule.
Remarks	One or more remarks configured for the selected ACL and associated with the rule during rule creation. To delete a remark associated with the rule, click the - (minus) button preceding remark. You must confirm the action before the rule associated remark is removed.
ACL Remarks	Lists the configured remarks for the selected ACL. All remarks present in this table are applied to the next rule created with the Add Rule button. Use the buttons available in the ACL Remarks table to perform the following tasks: <ul style="list-style-type: none"> To add a remark, click the + (plus) button and enter the remark to add. To delete a remark from the list, click the - (minus) button associated with the entry to remove. You must confirm the action before the entry is removed.

After you click **Add Rule**, the **Add Access Control List Rule** window opens and allows you to add a rule to the ACL that was selected from the **ACL Identifier** field. The fields available in the window depend on the ACL type. The following information describes the fields in this window. The Match Criteria tables that apply to IPv4 ACLs, IPv6 ACLs, and MAC ACLs are described separately.

Table 300: Add Access Control List Rule Fields

Field	Description
Match Criteria (IPv4 ACLs)	The fields in this section specify the criteria to use to determine whether an IP packet matches the rule. The fields described below apply to IPv4 Standard, IPv4 Extended, and IPv4 Named ACLs unless otherwise noted.
Every	When this option is selected, all packets will match the rule and will be either permitted or denied. This option is exclusive to all other match criteria, so if Every is selected, no other match criteria can be configured. To configure specific match criteria, this option must be clear.
Protocol	(IPv4 Extended and IPv4 Named ACLs) The IANA-assigned protocol number to match within the IP packet. You can also specify one of the following keywords: EIGRP, GRE, ICMP, IGMP, , IP, IPINIP, OSPF, PIM, TCP, or UDP.
Fragments	(IPv4 Extended and IPv4 Named ACLs) IP ACL rule to match on fragmented IP packets.
Source IP Address / Wildcard Mask	The source port IP address in the packet and source IP wildcard mask (in the second field) to compare to the IP address in a packet header. Wild card masks determines which bits in the IP address are used and which bits are ignored. A wild card mask of 255.255.255.255 indicates that no bit is important. A wildcard of 0.0.0.0 indicates that all of the bits are important. Wildcard masking for ACLs operates differently from a subnet mask. A wildcard mask is in essence the inverse of a subnet mask. With a subnet mask, the mask has ones (1's) in the bit positions that are used for the network address, and has zeros for the bit positions that are not used. In contrast, a wildcard mask has zeros in a bit position that must be checked. A '1' in a bit position of the ACL mask indicates the corresponding bit can be ignored. This field is required when you configure a source IP address.

Table 300: Add Access Control List Rule Fields (continued)

Field	Description
Source L4 Port	<p>(IPv4 Extended and IPv4 Named ACLs) The TCP/UDP source port to match in the packet header. The Source L4 Port and Destination L4 port are configurable only if protocol is either TCP or UDP. Equal to, Not Equal to, Greater than, and Less than options are available.</p> <ul style="list-style-type: none"> • For TCP protocol: BGP, Domain, Echo, FTP, FTP-Data, HTTP, SMTP, Telnet, WWW, POP2, or POP3 • For UDP protocol: Domain, Echo, NTP, RIP, SNMP, TFTP, Time, or WHO
Destination IP Address / Wildcard Mask	<p>The destination port IP address in the packet and destination IP wildcard mask (in the second field) to compare to the IP address in a packet header. Wild card masks determines which bits in the IP address are used and which bits are ignored. A wild card mask of 255.255.255.255 indicates that no bit is important. A wildcard of 0.0.0.0 indicates that all of the bits are important. Wildcard masking for ACLs operates differently from a subnet mask. A wildcard mask is in essence the inverse of a subnet mask. With a subnet mask, the mask has ones (1's) in the bit positions that are used for the network address, and has zeros for the bit positions that are not used. In contrast, a wildcard mask has zeros in a bit position that must be checked. A 1 in a bit position of the ACL mask indicates the corresponding bit can be ignored. This field is required when you configure a destination IP address.</p>
Destination L4 Port	<p>(IPv4 Extended and IPv4 Named ACLs) The TCP/UDP destination port to match in the packet header. The Source L4 Port and Destination L4 port are configurable only if protocol is either TCP or UDP. Equal to, Not Equal to, Greater than, and Less than options are available.</p> <ul style="list-style-type: none"> • For TCP protocol: BGP, Domain, Echo, FTP, FTP-Data, HTTP, SMTP, Telnet, WWW, POP2, or POP3 • For UDP protocol: Domain, Echo, NTP, RIP, SNMP, TFTP, Time, or WHO
TTL Field Value	<p>(IPv4 Extended and IPv4 Named ACLs) IP ACL rule to match on the specified Time-to-Live (TTL) field value.</p>
IGMP Type	<p>(IPv4 Extended and IPv4 Named ACLs) IP ACL rule to match on the specified IGMP message type. This option is available only if the protocol is IGMP.</p>
ICMP Type	<p>(IPv4 Extended and IPv4 Named ACLs) IP ACL rule to match on the specified ICMP message type. This option is available only if the protocol is ICMP.</p>
ICMP Code	<p>(IPv4 Extended and IPv4 Named ACLs) IP ACL rule to match on the specified ICMP message code. This option is available only if the protocol is ICMP.</p>
ICMP Message	<p>(IPv4 Extended and IPv4 Named ACLs) IP ACL rule to match on the ICMP message type and code. Specify one of the following supported ICMP messages: Echo, Echo-Reply, Host-Redirect, Mobile-Redirect, Net-Redirect, Net-Unreachable, Redirect, Packet-Too-Big, Port-Unreachable, Source-Quench, Router-Solicitation, Router-Advertisement, Time-Exceeded, TTL-Exceeded, and Unreachable. This option is available only if the protocol is ICMP.</p>
TCP Flags	<p>(IPv4 Extended and IPv4 Named ACLs) IP ACL rule to match on the TCP flags. When a + flag is specified, a match occurs if the flag is set in the TCP header. When a - flag is specified, a match occurs if the flag is not set in the TCP header. When Established is specified, a match occurs if either RST or ACK bits are set in the TCP header. This option is available only if the protocol is TCP.</p>

Table 300: Add Access Control List Rule Fields (continued)

Field	Description
Service Type	<p>(IPv4 Extended and IPv4 Named ACLs) The service type to match in the IP header. The options in this menu are alternative ways of specifying a match condition for the same Service Type field in the IP header, but each service type uses a different user notation. After you select the service type, specify the value for the service type in the appropriate field. Only the field associated with the selected service type can be configured. The services types are as follows:</p> <ul style="list-style-type: none"> • IP DSCP: Matches the packet IP DiffServ Code Point (DSCP) value to the rule. The DSCP value is defined as the high-order six bits of the Service Type octet in the IP header. • IP Precedence: Matches the IP Precedence value to the rule. The IP Precedence field in a packet is defined as the high-order three bits of the Service Type octet in the IP header. • IP TOS Bits: Matches on the Type of Service (TOS) bits in the IP header. The IP TOS field in a packet is defined as all eight bits of the Service Type octet in the IP header. For example, to check for an IP TOS value having bits 7 and 5 set and bit 1 clear, where bit 7 is most significant, use a TOS Bits value of 0xA0 and a TOS Mask of 0xFF. • TOS Bits: Requires the bits in a packet's TOS field to match the two-digit hexadecimal number entered in this field. • TOS Mask: The bit positions that are used for comparison against the IP TOS field in a packet.
Time Range Name	The name of the time range that will impose a time limitation on the ACL rule. If a time range with the specified name does not exist, and the ACL containing this ACL rule is associated with an interface, the ACL rule is applied immediately. If a time range with specified name exists, and the ACL containing this ACL rule is associated with an interface, the ACL rule is applied when the time-range with specified name becomes active. The ACL rule is removed when the time-range with specified name becomes inactive.
Committed Rate / Burst Size	The allowed transmission rate for packets on the interface (Committed Rate), and the number of bytes allowed in a temporary traffic burst (Burst Rate).
Match Criteria (IPv6 ACLs)	The fields in this section specify the criteria to use to determine whether an IP packet matches the rule. The fields described below apply to IPv6 ACLs.
Every	When this option is selected, all packets will match the rule and will be either permitted or denied. This option is exclusive to all other match criteria, so if Every is selected, no other match criteria can be configured. To configure specific match criteria, this option must be clear.
Protocol	The IANA-assigned protocol number to match within the IP packet. You can also specify one of the following keywords: ICMP, IGMP, TCP, UDP, ICMPv6, or IP.
Fragments	IPv6 ACL rule to match on fragmented IP packets.
Source Prefix/Prefix Length	The IPv6 prefix combined with IPv6 prefix length of the network or host from which the packet is being sent.

Table 300: Add Access Control List Rule Fields (continued)

Field	Description
Source L4 Port	The TCP/UDP source port to match in the packet header. Select one of the following options: Equal, Not Equal, Less Than, Greater Than, or Range and specify the port number or keyword. TCP port keywords include BGP, Domain, Echo, FTP, FTP Data, HTTP, SMTP, Telnet, WWW, POP2, and POP3. UDP port keywords include Domain, Echo, NTP, RIP, SNMP, TFTP, TIME, and WHO.
Destination Prefix/Prefix Length	The IPv6 prefix combined with the IPv6 prefix length to be compared to a packet's destination IPv6 address as a match criteria for the IPv6 ACL rule. To indicate a destination host, specify an IPv6 prefix length of 128.
Destination L4 Port	The TCP/UDP destination port to match in the packet header. Select one of the following options: Equal, Not Equal, Less Than, Greater Than, or Range and specify the port number or keyword. TCP port keywords include BGP, Domain, Echo, FTP, FTP Data, HTTP, SMTP, Telnet, WWW, POP2, and POP3. UDP port keywords include Domain, Echo, NTP, RIP, SNMP, TFTP, TIME, and WHO.
ICMP Type	IPv6 ACL rule to match on the specified ICMP message type. This option is available only if the protocol is ICMPv6.
ICMP Code	IPv6 ACL rule to match on the specified ICMP message code. This option is available only if the protocol is ICMPv6.
ICMP Message	IPv6 ACL rule to match on the ICMP message type and code. Specify one of the following supported ICMPv6 messages: Destination-Unreachable, Echo-Request, Echo-Reply, Header, Hop-Limit, MLD-Query, MLD-Reduction, MLD-Report, ND-NA, ND-NS, Next-Header, No-Admin, No-Route, Packet-Too-Big, Port-Unreachable, Router-Solicitation, Router-Advertisement, Router-Renumbering, Time-Exceeded, and Unreachable. This option is available only if the protocol is ICMPv6.
TCP Flags	IPv6 ACL rule to match on the TCP flags. When a + flag is specified, a match occurs if the flag is set in the TCP header. When a - flag is specified, a match occurs if the flag is not set in the TCP header. When Established is specified, a match occurs if either RST or ACK bits are set in the TCP header. This option is available only if the protocol is TCP.
Flow Label	A 20-bit number that is unique to an IPv6 packet, used by end stations to signify quality-of-service handling in routers.
IP DSCP	The IP DSCP value in the IPv6 packet to match to the rule. The DSCP value is defined as the high-order six bits of the Service Type octet in the IPv6 header.
Routing	IPv6 ACL rule to match on routed packets.
Match Criteria (MAC ACLs)	The fields in this section specify the criteria to use to determine whether an Ethernet frame matches the rule. The fields described below apply to MAC ACLs.
Every	When this option is selected, all packets will match the rule and will be either permitted or denied. This option is exclusive to all other match criteria, so if Every is selected, no other match criteria can be configured. To configure specific match criteria, this option must be clear.
CoS	The 802.1p user priority value to match within the Ethernet frame.

Table 300: Add Access Control List Rule Fields (continued)

Field	Description
Secondary CoS	The secondary 802.1p user priority value to match within the Ethernet frame.
Ethertype	The EtherType value to match in an Ethernet frame. Specify the number associated with the EtherType or specify one of the following keywords: AppleTalk, ARP, IBM SNA, IPv4, IPv6, IPX, MPLS, Unicast, NETBIOS, NOVELL, PPPoE, or RARP.
Source MAC Address / Mask	The MAC address to match to an Ethernet frame's source port MAC address. If desired, enter the MAC mask associated with the source MAC to match. The MAC address mask specifies which bits in the source MAC to compare against an Ethernet frame. Use Fs and zeros in the MAC mask, which is in a wildcard format. An F means that the bit is not checked, and a zero in a bit position means that the data must equal the value given for that bit. For example, if the MAC address is aa_bb_cc_dd_ee_ff, and the mask is 00_00_ff_ff_ff_ff, all MAC addresses with aa_bb_xx_xx_xx_xx result in a match (where x is any hexadecimal number).
Destination MAC Address / Mask	The MAC address to match to an Ethernet frame's destination port MAC address. If desired, enter the MAC Mask associated with the destination MAC to match. The MAC address mask specifies which bits in the destination MAC to compare against an Ethernet frame. Use F's and zeros in the MAC mask, which is in a wildcard format. An F means that the bit is not checked, and a zero in a bit position means that the data must equal the value given for that bit. For example, if the MAC address is aa_bb_cc_dd_ee_ff, and the mask is 00_00_ff_ff_ff_ff, all MAC addresses with aa_bb_xx_xx_xx_xx result in a match (where x is any hexadecimal number).
VLAN	The VLAN ID to match within the Ethernet frame.
Secondary VLAN	The secondary VLAN ID to match within the Ethernet frame.
Rule Attributes	The fields in this section provide information about the actions to take on a frame or packet that matches the rule criteria. The attributes specify actions other than the basic Permit or Deny actions.
Assign Queue	The number that identifies the hardware egress queue that will handle all packets matching this rule.
Interface	The interface to use for the action: <ul style="list-style-type: none"> • Redirect: Allows traffic that matches a rule to be redirected to the selected interface instead of being processed on the original port. The redirect function and mirror function are mutually exclusive. • Mirror: Provides the ability to mirror traffic that matches a rule to the selected interface. Mirroring is similar to the redirect function, except that in flow-based mirroring a copy of the permitted traffic is delivered to the mirror interface while the packet itself is forwarded normally through the device.
Log	When this option is selected, logging is enabled for this ACL rule (subject to resource availability in the device). If the Access List Trap Flag is also enabled, this will cause periodic traps to be generated indicating the number of times this rule went into effect during the current report interval. A fixed 5 minute report interval is used for the entire system. A trap is not issued if the ACL rule hit count is zero for the current interval.

Table 300: Add Access Control List Rule Fields (continued)

Field	Description
Time Range Name	The name of the time range that will impose a time limitation on the ACL rule. If a time range with the specified name does not exist, and the ACL containing this ACL rule is associated with an interface, the ACL rule is applied immediately. If a time range with specified name exists, and the ACL containing this ACL rule is associated with an interface, the ACL rule is applied when the time-range with specified name becomes active. The ACL rule is removed when the time-range with specified name becomes inactive.
Committed Rate / Burst Size	The allowed transmission rate for frames on the interface (Committed Rate), and the number of bytes allowed in a temporary traffic burst (Burst Rate).
After you click the Resequence Rules button, the Resequence ACL Rules window opens and allows you to resequence rules of the ACL selected from the ACL Identifier field. The following information describes the fields in this window.	
Sequence Start	The starting sequence number for resequencing the existing rules.
Sequence Step	The increment of sequence numbers for resequencing the existing rules.

Click **Refresh** to update the information on the screen.

After you click the + (plus) button next to **ACL Remarks**, the **Add ACL Remark** window opens and allows you to add a remark.

Access Control List Interface Summary

Use the **Access Control List Interface Summary** page to associate an ACL, or multiple ACLs, with one or more interfaces on the device. When an ACL is associated with an interface, traffic on the port is checked against the rules defined within the ACL until a match is found. If the traffic does not match any rules within an ACL, it is dropped because of the implicit deny all rule at the end of each ACL.

To access this page, click **QoS > Access Control Lists > Interfaces** in the navigation menu.

Use the buttons to perform the following tasks:

- To apply an ACL to an interface, click **Add** and configure the settings in the available fields.
- To remove the association between an interface and an ACL, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.

Table 301: Access Control List Interface Summary Fields

Field	Description
Interface	The interface that has an associated ACL.
Direction	Whether the packet is checked against the rules in an ACL when it is received on an interface (Inbound) or after it has been received, routed, and is ready to exit an interface (Outbound).
Sequence Number	The order the ACL is applied to traffic on the interface relative to other ACLs associated with the interface in the same direction. When multiple ACLs are applied to the same interface in the same direction, the ACL with the lowest sequence number is applied first, and the other ACLs are applied in ascending numerical order.

Table 301: Access Control List Interface Summary Fields (continued)

Field	Description
ACL Type	The type of ACL, which is either IPv4, IPv6, or MAC.
ACL Identifier	The name or number that identifies the ACL. When applying an ACL to an interface, the ACL Identifier menu includes only the ACLs within the selected ACL Type.

Access Control List VLAN Summary

Use the **Access Control List VLAN Summary** page to associate an [ACL](#), or multiple ACLs, with one or more VLANs on the device.



Note

You can also associate an ACL with a VLAN routing interface.

To access this page, click **QoS > Access Control Lists > VLANs** in the navigation menu.

Use the buttons to perform the following tasks:

- To associate an ACL with a VLAN, click **Add** and configure the settings in the available fields.
- To remove the association between a VLAN and an ACL, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.

Table 302: Access Control List VLAN Summary Fields

Field	Description
VLAN ID	The ID of the VLAN associated with the rest of the data in the row. When associating a VLAN with an ACL, use this field to select the desired VLAN.
Direction	Whether the packet is checked against the rules in an ACL when it is received on a VLAN (Inbound) or after it has been received, routed, and is ready to exit a VLAN (Outbound).
Sequence Number	The order the ACL is applied to traffic on the VLAN relative to other ACLs associated with the VLAN in the same direction. When multiple ACLs are applied to the same VLAN in the same direction, the ACL with the lowest sequence number is applied first, and the other ACLs are applied in ascending numerical order.
ACL Type	The type of ACL, which is either IPv4, IPv6, or MAC.
ACL Identifier	The name or number that identifies the ACL. The permitted identifier depends on the ACL type. Standard and Extended IPv4 ACLs use numbers within a set range, and Named IPv4, IPv6, and MAC ACLs use alphanumeric characters.

Access Control List Control Plane Configuration

Use the **Access Control List Control Plane Configuration** page to define controlled management access to the device. A control plane [ACL](#) enables you to determine which addresses or protocols are allowed to access the management interface on the device. The control plane ACLs are applied to management access through the in-band (production network) ports only. Inbound traffic on the CPU port is checked

against the rules defined within the ACL until a match is found. If the traffic does not match any rules within an ACL, it is dropped because of the implicit deny all rule at the end of each ACL.

To access this page, click **QoS > Access Control Lists > Control Plane** in the navigation menu.

Use the buttons to perform the following tasks:

- To apply an ACL to the CPU interface, click **Add** and configure the settings in the available fields.
- To remove the association between the CPU interface and an ACL, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.

Table 303: Access Control List Control Plane Configuration Fields

Field	Description
ACL Identifier	The name or number that identifies the ACL.
ACL Type	<p>The type of ACL. The ACL type determines the criteria that can be used to match packets. The type also determines which attributes can be applied to matching traffic. IPv4 ACLs classify Layer 3 and Layer 4 IPv4 traffic, IPv6 ACLs classify Layer 3 and Layer 4 IPv6 traffic, and MAC ACLs classify Layer 2 traffic. The ACL types are as follows:</p> <ul style="list-style-type: none"> • IPv4 Standard: Match criteria is based on the source address of IPv4 packets. • IPv4 Extended: Match criteria can be based on the source and destination addresses, source and destination Layer 4 ports, and protocol type of IPv4 packets. • IPv4 Named: Match criteria is the same as IPv4 Extended ACLs, but the ACL ID can be an alphanumeric name instead of a number. • IPv6 Named: Match criteria can be based on information including the source and destination IPv6 addresses, source and destination Layer 4 ports, and protocol type within IPv6 packets. • Extended MAC: Match criteria can be based on the source and destination MAC addresses, 802.1p user priority, VLAN ID, and EtherType value within Ethernet frames.
Sequence Number	The order the ACL is applied to traffic on the interface relative to other ACLs associated with the interface in the same direction. When multiple ACLs are applied to the same interface in the same direction, the ACL with the lowest sequence number is applied first, and the other ACLs are applied in ascending numerical order.

IPv6 ACL Rules

The maximum number of IPv6 rules depends on the following factors (also refer to the 200 Series Scaling Parameters and Values for the maximum number of rules per device type):

- If both SRC IPv6 and DST IPv6 are part of the ACL rule, then the maximum number of rules is one quarter the possible number for that device type.
- If DSCP is part of the rule along with any other qualifier, then the maximum number of rules possible are one quarter the possible number for that device type.
- In all other cases, the maximum number of rules are equal to half the maximum possible for that device type or 1021, whichever is smaller.

Scenarios

The following scenarios are provided as an example. Assume that your hardware processor can accommodate a maximum of 1789 rules.

- Scenario #1: If the rules have both SRC IPv6 and DST IPv6, then maximum rules possible are $1789/4 = 447$.
- Scenario #2: If the rules have DSCP along with any other qualifier, then the maximum number of rules possible are $1789/4 = 447$.
- Scenario #3: In all the other cases, 894 rules can be accommodated.

Access Control List Statistics

Use the **Access Control List Statistics** page to display the statistical information about the packets forwarded or discarded by the port that matches the configured rules within an *ACL*. Each ACL rule is configured to match one or more aspects of traffic on the network. When a packet matches the conditions in a rule, the counter associated with the rule gets incremented, until it reaches the rollover value of the counter. ACL counters do not interact with DiffServ policies or Policy-based Routing counters.

To access this page, click **QoS > Access Control Lists > Statistics** in the navigation menu.

Use the buttons to perform the following tasks:

- To clear the hit count for one or more configured rules within an ACL, select the rule entry and click **Clear Rule Counter**. You must confirm the action before the hit count is cleared for the selected rule(s).
- To clear the hit count for an ACL, select the ACL ID from the ACL Identifier menu and click **Clear ACL Counters**. You must confirm the action before the hit count is cleared for the selected ACL.
- To clear the hit count for an ACL type, select the type from the ACL Type menu and select **All** from the ACL Identifier menu and then click **Clear ACL Counters**. You must confirm the action before the hit count is cleared for the selected ACL type.

Table 304: Access Control List Statistics Fields

Field	Description
ACL Type	<p>The type of ACL. The ACL type determines the criteria that can be used to match packets. The type also determines which attributes can be applied to matching traffic. IPv4 ACLs classify Layer 3 and Layer 4 IPv4 traffic, IPv6 ACLs classify Layer 3 and Layer 4 IPv6 traffic, and MAC ACLs classify Layer 2 traffic. The ACL types are as follows:</p> <ul style="list-style-type: none"> • IPv4 Standard: Match criteria is based on the source address of the IPv4 packets. • IPv4 Extended: Match criteria can be based on the source and destination addresses, source and destination Layer 4 ports, and protocol type of the IPv4 packets. • IPv4 Named: Match criteria is the same as IPv4 Extended ACLs, but the ACL ID can be an alphanumeric name instead of a number. • IPv6 Named: Match criteria can be based on information including the source and destination IPv6 addresses, source and destination Layer 4 ports, and protocol type within the IPv6 packets. • Extended MAC: Match criteria can be based on the source and destination MAC addresses, 802.1p user priority, VLAN ID, and EtherType value within the Ethernet frames.
ACL Identifier	A list of ACL IDs that exist on the system for a given ACL type. To view the rule(s) within an ACL, you must select the ID of the ACL from the list. The ACL rules are not displayed when option All is selected. Option All lets you clear the hit count for an ACL type.
Sequence Number	The number that indicates the position of a rule within the ACL.
Action	<p>The action to take when a packet or frame matches the criteria in the rule:</p> <ul style="list-style-type: none"> • Permit: The packet or frame is forwarded. • Deny: The packet or frame is dropped.
Match Conditions	The criteria used to determine whether a packet or frame matches the ACL rule.
Rule Attributes	Each action — beyond the basic Permit and Deny actions — to perform on the traffic that matches the rule.
Hit Count	Indicates the number of packets that match the configured rule in an ACL. If a rule is configured without rate limit, then the hit count is the number of matched packets forwarded or discarded by the port. If a rule is configured with rate limit, then if the sent traffic rate exceeds the configured rate, the hit count displays the matched packet count equal to the sent rate, despite packets getting dropped beyond the configured limit. If the sent traffic rate is less than the configured rate, the hit count displays only the matched packet count.

Configuring Auto VoIP

Voice over Internet Protocol (VoIP) allows you to make telephone calls using a computer network over a data network like the Internet. With the increased prominence of delay-sensitive applications (voice, video, and other multimedia applications) deployed in networks today, proper QoS configuration will ensure high-quality application performance. The Auto VoIP feature is intended to provide an easy classification mechanism for voice packets so that they can be prioritized above data packets in order to provide better QoS.

The Auto-VoIP feature explicitly matches VoIP streams in Ethernet switches and provides them with a better class of service than ordinary traffic. If you enable the Auto-VoIP feature on an interface, the interface scans incoming traffic for the following call-control protocols:

- Session Initiation Protocol (SIP)
- H.323
- Skinny Client Control Protocol (SCCP)

When a call-control protocol is detected the switch assigns the traffic in that session to the highest CoS (Class of Service) queue, which is generally used for time-sensitive traffic.

Auto VoIP Global Configuration

Use the **Auto VoIP Global Configuration** page to configure the VLAN ID for the Auto VoIP VLAN or to reset the current Auto VoIP VLAN ID to the default value. Voice over Internet Protocol (VoIP) enables telephone calls over a data network. Because voice traffic is typically more time-sensitive than data traffic, the Auto VoIP feature helps provide a classification mechanism for voice packets so that they can be prioritized above data packets in order to provide better QoS.

With the Auto VoIP feature, voice prioritization is provided based on call-control protocols (SIP, SCCP, H.323) and/or OUI bits. When the device identifies voice traffic, it is placed in the VLAN specified on this page. The Auto VoIP feature does not rely on LLDP-MED support in connected devices.

To access this page, click **QoS > Auto VoIP > Global** in the navigation menu.

Table 305: Auto VoIP Global Configuration Fields

Field	Description
Auto VoIP VLAN	The VLAN used to segregate VoIP traffic from other non-voice traffic.
Reset (Button)	Click this button to reset the voice VLAN to the default value.

OUI Table Summary

Use the **OUI Table Summary** page to add and remove Organizationally Unique Identifiers (OUIs) from the OUI database the device maintains. Device hardware manufacturers can include an OUI in a network adapter to help identify the device. The OUI is a unique 24-bit number assigned by the IEEE registration authority. Several default OUIs have been preconfigured in the OUI database on the device.

To access this page, click **QoS > Auto VoIP > OUI Table** in the navigation menu.

Use the buttons to perform the following tasks:

- To add an OUI, click **Add** and specify an OUI and its description in the available fields.
- To remove one or more configured OUIs, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.

Table 306: OUI Table Summary Fields

Field	Description
Telephony OUI	The unique OUI that identifies the device manufacturer or vendor. The OUI is specified in three octet values (each octet is represented as two hexadecimal digits) separated by colons.
Status	Identifies whether the OUI is preconfigured on the system (Default) or added by a user (Configured).
Description	Identifies the manufacturer or vendor associated with the OUI.

OUI Based Auto VoIP

Use the **OUI Based Auto VoIP** page to configure the Organizationally Unique Identifier (OUI) based Auto VoIP priority and to enable or disable the Auto VoIP mode on the interfaces.

To access this page, click **QoS > Auto VoIP > OUI Based Auto VoIP** in the navigation menu.

Use the buttons to perform the following tasks:

- To configure the settings for one or more interfaces, select each entry to modify and click **Edit**.
- To apply the same settings to all interfaces, click **Edit All**.

Table 307: OUI Based Auto VoIP Fields

Field	Description
Auto VoIP VLAN	The VLAN used to segregate VoIP traffic from other non-voice traffic. All VoIP traffic that matches a value in the known OUI list gets assigned to this VoIP VLAN.
Priority	The 802.1p priority used for traffic that matches a value in the known OUI list. If the Auto VoIP mode is enabled and the interface detects an OUI match, the device assigns the traffic in that session to the traffic class mapped to this priority value. Traffic classes with a higher value are generally used for time-sensitive traffic.
Interface	The interface associated with the rest of the data in the row. When editing Auto VoIP settings on one or more interfaces, this field identifies the interface(s) being configured.
Auto VoIP Mode	The administrative mode of OUI-based Auto VoIP on the interface.
Operational Status	The operational status of an interface. To be up, an interface must be administratively enabled and have a link.

Protocol Based Auto VoIP

Use the **Protocol Based Auto VoIP** page to configure the protocol-based Auto VoIP priority settings and to enable or disable the protocol-based Auto VoIP mode on the interfaces.

To access this page, click **QoS > Auto VoIP > Protocol Based Auto VoIP** in the navigation menu.

Use the buttons to perform the following tasks:

- To configure the settings for one or more interfaces, select each entry to modify and click **Edit**.
- To apply the same settings to all interfaces, click **Edit All**.
- If you change any of the settings on the page, click **Submit** to send the updated configuration to the switch. Configuration changes take effect immediately. These changes will not be retained across a power cycle unless a Save is performed.
- Click **Refresh** to update the page with the most current data from the switch.

Table 308: Protocol Based Auto VoIP Fields

Field	Description
Auto VoIP VLAN	The VLAN used to segregate VoIP traffic from other non-voice traffic. All VoIP traffic in a session identified by the call-control protocol gets assigned to this VoIP VLAN.
Prioritization Type	The method used to prioritize VoIP traffic when a call-control protocol is detected, which is one of the following: <ul style="list-style-type: none"> • Remark: Remark the voice traffic with the specified 802.1p priority value at the ingress interface. • Traffic Class: Assign VoIP traffic to the specified traffic class when egressing the interface.
802.1p Priority	The 802.1p priority used for protocol-based VoIP traffic. This field can be configured if the Prioritization Type is 802.1p Priority. If the Auto VoIP mode is enabled and the interface detects a call-control protocol, the device marks traffic in that session with the specified 802.1p priority value to ensure voice traffic always gets the highest priority throughout the network path. Egress tagging must be administratively enabled on the appropriate uplink port to carry the remarked priority at the egress port.
Traffic Class	The traffic class used for protocol-based VoIP traffic. This field can be configured if the Prioritization Type is Traffic Class. If the Auto VoIP mode is enabled and the interface detects a call-control protocol, the device assigns the traffic in that session to the configured CoS queue. Traffic classes with a higher value are generally used for time-sensitive traffic. The CoS queue associated with the specified traffic class should be configured with the appropriate bandwidth allocation to allow priority treatment for VoIP traffic.
Interface	The interface associated with the rest of the data in the row. When editing Auto VoIP settings on one or more interfaces, this field identifies the interface(s) being configured.
Auto VoIP Mode	The administrative mode of the Auto VoIP feature on the interface: <ul style="list-style-type: none"> • Enable: The interface scans incoming traffic for the following call-control protocols: <ul style="list-style-type: none"> • Session Initiation Protocol (SIP) • H.323 • Skinny Client Control Protocol (SCCP) • Disable: The interface does not use the Auto VoIP feature to scan for call-control protocols.
Operational Status	The operational status of an interface. To be up, an interface must be administratively enabled and have a link.

Configuring Class of Service

The CoS queueing feature lets you directly configure certain aspects of switch queueing. This provides the desired QoS behavior for different types of network traffic when the complexities of DiffServ are not required. The priority of a packet arriving at an interface can be used to steer the packet to the appropriate outbound CoS queue through a mapping table. CoS queue characteristics that affect queue mapping, such as minimum guaranteed bandwidth, transmission rate shaping, etc., are user-configurable at the queue (or port) level.

Seven queues per port are supported. Although the hardware supports eight queues, one queue is always reserved for internal use by the stacking subsystem.

IP DSCP Mapping Configuration

Use the **IP DSCP Mapping Configuration** page to map an IP DSCP value to an internal traffic class.

To access this page, click **QoS > Class of Service > IP DSCP** in the navigation menu.

Table 309: IP DSCP Mapping Configuration Fields

Field	Description
Interface	The menu contains all CoS configurable interfaces. The only option is Global, which means that the IP DSCP mapping configuration applies to all interfaces and cannot be applied on a per-interface basis.
IP DSCP Values	Lists the IP DSCP values to which you can map an internal traffic class. The values range from 0 to 63.
Traffic Class	The traffic class is the hardware queue for a port. Higher traffic class values indicate a higher queue position. Before traffic in a lower queue is sent, it must wait for traffic in higher queues to be sent. Valid range is 0 to 6.

If you make changes to the page, click **Submit** to apply the changes to the system. Click **Restore Defaults** to reset all interfaces to the default trust value.

Interface Configuration

Use the **Interface Configuration** page to apply an interface shaping rate to all ports or to a specific port.

To access this page, click **QoS > Class of Service > Interface** in the navigation menu.

Table 310: Interface Configuration Fields

Field	Description
Interface	Selects the CoS configurable interface to be affected by the Interface Shaping Rate. Select Global to apply a rate to all interfaces. Select an individual port to override the global setting.
Interface Shaping Rate	Sets the limit on how much traffic can leave a port. The limit on maximum transmission bandwidth has the effect of smoothing temporary traffic bursts over time so that the transmitted traffic rate is bounded. The specified value represents a percentage of the maximum negotiated bandwidth. The default value is zero (0). Valid values are 0 to 100, in increments of 1. A value of 0 means the maximum is unlimited.
WRED Decay Exponent	Specifies the decay exponent value used with the WRED average queue length calculation algorithm. Default value is 9. Valid Range is (0 to 15).

If you make changes to the page, click **Submit** to apply the changes to the system. Click Restore Defaults to reset all interfaces to the default trust value.

Interface Queue Configuration

Use the **Interface Queue Configuration** page to define what a particular queue does by configuring switch egress queues. User-configurable parameters control the amount of bandwidth used by the queue, the queue depth during times of congestion, and the scheduling of packet transmission from the set of all queues on a port. Each port has its own CoS queue-related configuration.

The configuration process is simplified by allowing each CoS queue parameter to be configured globally or per-port. A global configuration change is automatically applied to all ports in the system.

To access this page, click **QoS > Class of Service > Queue** in the navigation menu.

Table 311: Interface Queue Configuration Fields

Field	Description
Interface	Specifies the interface (physical, LAG (Link Aggregation Group) , or Global) to configure.
Minimum Bandwidth Allocated	Shows the sum of individual Minimum Bandwidth values for all queues in the interface. The sum cannot exceed the defined maximum of 100. This value is considered while configuring the Minimum Bandwidth for a queue in the selected interface.
Queue ID	Use the menu to select the queue per interface to be configured.
Minimum Bandwidth	Specify the minimum guaranteed bandwidth allocated to the selected queue on the interface. Setting this value higher than its corresponding Maximum Bandwidth automatically increases the maximum to the same value. The default value is 0. The valid range is 0 to 100, in increments of 1. The value zero (0) means no guaranteed minimum. The sum of individual Minimum Bandwidth values for all queues in the selected interface cannot exceed defined maximum 100.

Table 311: Interface Queue Configuration Fields (continued)

Field	Description
Scheduler Type	<p>Selects the type of queue processing from the drop-down menu. Options are Weighted and Strict. Defining on a per-queue basis allows the user to create the desired service characteristics for different types of traffic.</p> <ul style="list-style-type: none"> • Weighted: Weighted round robin associates a weight to each queue. This is the default. • Strict: Strict priority services traffic with the highest priority on a queue first
Queue Management Type	<p>Displays the type of queue depth management techniques used for all queues on this interface. This is only used if the device supports independent settings per-queue. Queue Management Type can only be Taildrop. The default value is Taildrop. All packets on a queue are safe until congestion occurs. At this point, any additional packets queued are dropped.</p>

- If you make changes to the page, click **Submit** to apply the changes to the system.
- Click **Restore Defaults for all Queues** to reset the settings for the selected interface.
- To reset the defaults for all interfaces, select Global from the **Slot/Port** menu before you click the button.

Configuring Diffserv

Use this page to configure the administrative mode of Differentiated Services (DiffServ) support on the device and to view the current and maximum number of entries in each of the main DiffServ private MIB tables. DiffServ allows traffic to be classified into streams and given certain *QoS* treatment in accordance with defined per-hop behaviors.

Packets are classified and processed based on defined criteria. The classification criteria is defined by a class. The processing is defined by a policy's attributes. Policy attributes may be defined on a per-class instance basis, and it is these attributes that are applied when a match occurs. A policy can contain multiples classes. When the policy is active, the actions taken depend on which class matches the packet.

Diffserv Global Configuration and Status

Use the **Diffserv Global Configuration and Status** page to configure the Global DiffServ settings on the device.

To access this page, click **QoS > Diffserv > Global** in the navigation menu.

Table 312: Diffserv Global Configuration and Status Fields

Field	Description
Diffserv Admin Mode	The administrative mode of DiffServ on the device. While disabled, the DiffServ configuration is retained and can be changed, but it is not active. While enabled, Differentiated Services are active.
MIB Table	The information in this table displays the number of entries (rows) that are currently in each of the main DiffServ private MIB tables and the maximum number of rows that can exist in each table.

Table 312: Diffserv Global Configuration and Status Fields (continued)

Field	Description
Class Table	The current and maximum number of classifier entries in the table. DiffServ classifiers differentiate among traffic types.
Class Rule Table	The current and maximum number of class rule entries in the table. Class rules specify the match criteria that belong to a class definition.
Policy Table	The current and maximum number of policy entries in the table. The policy determines the traffic conditioning or service provisioning actions applied to a traffic class.
Policy Instance Table	The current and maximum number of policy-class instance entries in the table. A policy-class instance is a policy that is associated with an existing DiffServ class.
Policy Attribute Table	The current and maximum number of policy attribute entries in the table. A policy attribute entry attaches various policy attributes to a policy-class instance.
Service Table	The current and maximum number of service entries in the table. A service entry associates a DiffServ policy with an interface and inbound or outbound direction.

- If you make changes to the page, click **Submit** to apply the changes to the system. Click Restore Defaults to reset all interfaces to the default trust value.
- Click **Refresh** to update the page with the most current data from the switch.

Diffserv Class Summary

Use the **Diffserv Class Summary** page to create or remove DiffServ classes and to view summary information about the classes that exist on the device. Creating a class is the first step in using DiffServ to provide QoS. After a class is created, you can define the match criteria for the class.

To access this page, click **QoS > Diffserv > Class Summary** in the navigation menu.

Use the buttons to perform the following tasks:

- To add a DiffServ class, click **Add**.
- To change the name of an existing class, select the entry to modify and click **Rename**.
- To remove one or more configured classes, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.

Table 313: Diffserv Class Summary Fields

Field	Description
Name	The name of the DiffServ class. When adding a new class or renaming an existing class, the name of the class is specified in the Class field of the dialog window.
Type	The class type, which is one of the following: <ul style="list-style-type: none"> • All: All the various match criteria defined for the class should be satisfied for a packet match. All signifies the logical AND of all the match criteria. • Any: Any of various match criteria defined for the class can be satisfied for a packet match. • ACL: The criteria defined in the associated <u>ACL</u> is used to match packets. When adding a class in the Add Class window, the ACL (IP) option is available only if at least one IP ACL exists on the device. Similarly, the ACL (MAC) option is available only if at least one MAC ACL exists on the device.
Protocol	The Layer 3 protocol to use for filtering class types, which is either IPv4 or IPv6.
Match Criteria	The criteria used to match packets.
After you click Add , the Add Class window opens. The following information describes the additional field that appears in this window.	
ACL	If the selected Type is ACL (IP) or ACL (MAC), use this menu select the ACL to use to match packets.

Click **Refresh** to update the page with the most current data from the switch.

Diffserv Class Configuration

Use the **Diffserv Class Configuration** page to define the criteria to associate with a DiffServ class. As packets are received or transmitted, these DiffServ classes are used to classify and prioritize packets. Each class can contain multiple match criteria.

After you select the class to configure from the **Class** menu, use the buttons to perform the following tasks:

- To define criteria for matching packets within a class, click **Add Match Criteria**. Once you add a match criteria entry to a class, you cannot edit or remove the entry. However, you can add more match criteria entries to a class until the maximum number of entries has been reached for the class. This button is not available if the class type is ACL because the match criteria are defined by the ACL rules.
- To remove the associated reference class from the selected class, click Remove Reference Class. Note that unless the reference class is the last entry in the list of match criteria, the Reference Class match type remains in the list as a placeholder, but the associated value is N/A, and the previously referenced class is removed.

To access this page, click **QoS > Diffserv > Class Configuration** in the navigation menu.

Table 314: Diffserv Class Configuration Fields

Field	Description
Class	The name of the class. To configure match criteria for a class, select its name from the menu.
Type	The class type, which is one of the following: <ul style="list-style-type: none"> • All: All the various match criteria defined for the class should be satisfied for a packet match. All signifies the logical AND of all the match criteria. • Any: Any of various match criteria defined for the class can be satisfied for a packet match. • ACL: The criteria defined in the associated ACL is used to match packets. When adding a class in the Add Match Criteria window, the ACL (IP) option is available only if at least one IP ACL exists on the device. Similarly, the ACL (MAC) option is available only if at least one MAC ACL exists on the device.
L3 Protocol	The Layer 3 protocol to use for filtering class types, which is either IPv4 or IPv6.
Match Criteria	The type of match criteria defined for the selected class. If the Type is ACL, no information about the match criteria is available on this page.
Value	The configured value of the match criteria that corresponds to the match type.

After you click **Add Match Criteria**, the **Add Match Criteria** window opens and allows you to define the match criteria for the selected class. The window lists the match criteria that are available for the class. To add match criteria, select the checkbox for the criteria type. The fields to configure the match values appear after you select the match type. Each match criteria type can be used only once within a class. If a reference class includes the match criteria type, it cannot be used as an additional match type within the class, and the match criteria type cannot be selected or configured.

Each match type other than the Reference Class includes an option to match any value within the match criteria type except the configured value. This is the Exclude option, which indicates a logical NOT for a match criteria type.

Table 315: Add Match Criteria Fields

Field	Description
Any	Select this option to specify that all packets are considered to match the specified class. There is no need to configure additional match criteria if Any is selected because a match will occur on all packets.
Reference Class	Select this option to reference another class for criteria. The match criteria defined in the referenced class is as match criteria in addition to the match criteria you define for the selected class. After selecting this option, the classes that can be referenced are displayed. Select the class to reference. A class can reference at most one other class of the same type.
CoS	Select this option to require the Class of Service (CoS) value in an Ethernet frame header to match the specified CoS value.
Secondary Class of Service	Select this option to require the secondary CoS value in an Ethernet frame header to match the specified secondary CoS value.

Table 315: Add Match Criteria Fields (continued)

Field	Description
Ethertype	<p>Select this option to require the EtherType value in the Ethernet frame header to match the specified EtherType value. After you select this option, specify the EtherType value in one of the following two fields:</p> <ul style="list-style-type: none"> • Ethertype Keyword: The menu includes several common protocols that are mapped to their EtherType values. • Ethertype Value: This field accepts custom EtherType values.
VLAN	<p>Select this option to require a packet's VLAN ID to match a VLAN ID or a VLAN ID within a continuous range. If you configure a range, a match occurs if a packet's VLAN ID is the same as any VLAN ID within the range. After you select this option, use the following fields to configure the VLAN match criteria:</p> <ul style="list-style-type: none"> • VLAN ID Start: The VLAN ID to match or the VLAN ID with the lowest value within a range of VLANs. • VLAN ID End: The VLAN ID with the highest value within the range of VLANs. This field is not required if the match criteria is a single VLAN ID.
Secondary VLAN	<p>Select this option to require a packet's VLAN ID to match a secondary VLAN ID or a secondary VLAN ID within a continuous range. If you configure a range, a match occurs if a packet's secondary VLAN ID is the same as any secondary VLAN ID within the range. After you select this option, use the following fields to configure the secondary VLAN match criteria:</p> <ul style="list-style-type: none"> • Secondary VLAN ID Start: The secondary VLAN ID to match or the secondary VLAN ID with the lowest value within a range of VLANs. • Secondary VLAN ID End: The secondary VLAN ID with the highest value within the range of VLANs. This field is not required if the match criteria is a single VLAN ID.
Source MAC Address	<p>Select this option to require a packet's source MAC address to match the specified MAC address. After you select this option, use the following fields to configure the source MAC address match criteria:</p> <ul style="list-style-type: none"> • MAC Address: The source MAC address to match. • MAC Mask: The MAC mask, which specifies the bits in the source MAC address to compare against an Ethernet frame. Use F's and zeros to configure the MAC mask. An F means that the bit is checked, and a zero in a bit position means that the data is not significant. <p>For example, if the MAC address is aa:bb:cc:dd:ee:ff, and the mask is ff:ff:00:00:00:00, all MAC addresses with aa:bb:xx:xx:xx:xx result in a match (where x is any hexadecimal number). Note that this is not a wildcard mask, which ACLs use.</p>

Table 315: Add Match Criteria Fields (continued)

Field	Description
Destination MAC Address	<p>Select this option to require a packet's destination MAC address to match the specified MAC address. After you select this option, use the following fields to configure the destination MAC address match criteria:</p> <ul style="list-style-type: none"> • MAC Address: The destination MAC address to match. • MAC Mask: The MAC mask, which specifies the bits in the destination MAC address to compare against an Ethernet frame. Use Fs and zeros to configure the MAC mask. An F means that the bit is checked, and a zero in a bit position means that the data is not significant. <p>For example, if the MAC address is aa:bb:cc:dd:ee:ff, and the mask is ff:ff:00:00:00:00, all MAC addresses with aa:bb:xx:xx:xx:xx result in a match (where x is any hexadecimal number). Note that this is not a wildcard mask, which ACLs use.</p>
Source IPv6 Address	<p>Select this option to require the source IPv6 address in a packet header to match the specified values. After you select this option, use the following fields to configure the source IPv6 address match criteria:</p> <ul style="list-style-type: none"> • Source Prefix: The source IPv6 prefix to match. • Source Prefix Length: The IPv6 prefix length.
Destination IPv6 Address	<p>Select this option to require the destination IPv6 address in a packet header to match the specified values. After you select this option, use the following fields to configure the destination IPv6 address match criteria:</p> <ul style="list-style-type: none"> • Destination Prefix: The destination IPv6 prefix to match. • Destination Prefix Length: The IPv6 prefix length.
Source L4 Port	<p>Select this option to require a packet's TCP/UDP source port to match the specified port or the port number within a range of port numbers. If you configure a range, a match occurs if a packet's source port number is the same as any source port number within the range. After you select this option, use the following fields to configure a source port keyword, source port number, or source port range for the match criteria:</p> <ul style="list-style-type: none"> • Protocol: Select the desired L4 keyword from the list on which the match is based. If you select a keyword, the other source port configuration fields are not available. • Port End: A user-defined L4 source port number to match or the source port number with the lowest value within a range of ports. • Port Start: The source port with the highest value within the range of ports. This field is not required if the match criteria is a single port.

Table 315: Add Match Criteria Fields (continued)

Field	Description
Destination L4 Port	<p>Select this option to require a packet's TCP/UDP destination port to match the specified port or the port number within a range of port numbers. If you configure a range, a match occurs if a packet's destination port number is the same as any destination port number within the range. After you select this option, use the following fields to configure a destination port keyword, destination port number, or destination port range for the match criteria:</p> <ul style="list-style-type: none"> • Protocol: Select the desired L4 keyword from the list on which the match is based. If you select a keyword, the other destination port configuration fields are not available. • Port End: A user-defined L4 destination port number to match or the destination port number with the lowest value within a range of ports. • Port Start: The destination port with the highest value within the range of ports. This field is not required if the match criteria is a single port.
IP DSCP	<p>Select this option to require the packet's IP DiffServ Code Point (DSCP) value to match the specified value. The DSCP value is defined as the high-order six bits of the Service Type octet in the IP header. After you select this option, use one of the following fields to configure the IP DSCP match criteria:</p> <ul style="list-style-type: none"> • IP DSCP Keyword: The IP DSCP keyword code that corresponds to the IP DSCP value to match. If you select a keyword, you cannot configure an IP DSCP Value. • IP DSCP Value: The IP DSCP value to match.
IP Precedence	<p>Select this option to require the packet's IP Precedence value to match the number configured in the IP Precedence Value field. The IP Precedence field in a packet is defined as the high-order three bits of the Service Type octet in the IP header.</p>
IP TOS	<p>Select this option to require the packet's Type of Service (ToS) bits in the IP header to match the specified value. The IP ToS field in a packet is defined as all eight bits of the Service Type octet in the IP header. After you select this option, use the following fields to configure the ToS match criteria:</p> <ul style="list-style-type: none"> • IP TOS Bits: Enter a two-digit hexadecimal number to match the bits in a packet's ToS field. • IP TOS Mask: Specify the bit positions that are used for comparison against the IP ToS field in a packet.
Protocol	<p>Select this option to require a packet header's Layer 4 protocol to match the specified value. After you select this option, use one of the following fields to configure the protocol match criteria:</p> <ul style="list-style-type: none"> • Protocol: The L4 keyword that corresponds to value of the IANA protocol number to match. If you select a keyword, you cannot configure a Protocol Value. • Protocol Value: The IANA L4 protocol number value to match.
Flow Label	<p>Select this option to require an IPv6 packet's flow label to match the configured value. The flow label is a 20-bit number that is unique to an IPv6 packet, used by end stations to signify quality-of-service handling in routers.</p>

Click **Refresh** to update the page with the most current data from the switch.

Diffserv Policy Summary

Use the **Diffserv Policy Summary** page to create or remove DiffServ policies and to view summary information about the policies that exist on the device. A policy defines the **QoS** attributes for one or more traffic classes. A policy attribute identifies the action taken when a packet matches a class rule. A policy is applied to a packet when a class match within that policy is found.

To access this page, click **QoS > Diffserv > Policy Summary** in the navigation menu.

Use the buttons to perform the following tasks:

- To add a DiffServ policy, click **Add**.
- To change the name of an existing policy, select the entry to modify and click **Rename**.
- To remove one or more configured policies, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.

Table 316: Diffserv Policy Summary Fields

Field	Description
Name	The name of the DiffServ policy. When adding a new policy or renaming an existing policy, the name of the policy is specified in the Policy field of the dialog window.
Type	The traffic flow direction to which the policy is applied: <ul style="list-style-type: none"> • In: The policy is specific to inbound traffic. • Out: The policy is specific to outbound traffic direction.
Member Classes	The DiffServ class or classes that have been added to the policy.

Click **Refresh** to update the page with the most current data from the switch.

Diffserv Policy Configuration

Use the **Diffserv Policy Configuration** page to add or remove a DiffServ policy-class association and to configure the policy attributes. The policy attributes identify the action or actions taken when a packet matches a class rule.

After you select the policy to configure from the Policy menu, use the buttons to perform the following tasks:

- To add a class to the policy, click **Add Class**.
- To add attributes to a policy or to change the policy attributes, select the policy with the attributes to configure and click **Add Attribute**.
- To remove the most recently associated class from the selected policy, click **Remove Last Class**.

To access this page, click **QoS > Diffserv > Policy Configuration** in the navigation menu.

Table 317: Diffserv Policy Configuration Fields

Field	Description
Policy	The name of the policy. To add a class to the policy, remove a class from the policy, or configure the policy attributes, you must first select its name from the menu.
Type	The traffic flow direction to which the policy is applied.
Class	The DiffServ class or classes associated with the policy. The policy is applied to a packet when a class match within that policy-class is found.
Policy Attribute Details	The policy attribute types and their associated values that are configured for the policy.
Assign Queue	Select this option to assign matching packets to a traffic queue. Use the Queue ID Value field to select the queue to which the packets of this policy-class are assigned.
Drop	Select this option to drop packets that match the policy-class.
Mark CoS	Select this option to mark all packets in a traffic stream with the specified Class of Service (CoS) queue value. Use the Class of Service field to select the CoS value to mark in the priority field of the 802.1p header (the only tag in a single tagged packet or the first or outer 802.1Q tag of a double VLAN tagged packet). If the packet does not already contain this header, one is inserted.
Mark Secondary CoS	Select this option to mark all packets in a traffic stream with the specified secondary CoS queue number. Use the Class of Service field to select the CoS value to mark in the priority field of the 802.1p header in the secondary (inner) 802.1Q tag of a double VLAN tagged packet. If the packet does not already contain this header, one is inserted.
Mark IP DSCP	Select this option to mark all packets in the associated traffic stream with the specified IP DSCP value. After you select this option, use one of the following fields to configure the IP DSCP value to mark in packets that match the policy-class: <ul style="list-style-type: none"> • IP DSCP Keyword: The IP DSCP keyword code that corresponds to the IP DSCP value. If you select a keyword, you cannot configure an IP DSCP Value. • IP DSCP Value: The IP DSCP value.
Mark IP Precedence	Select this option to mark all packets in the associated traffic stream with the specified IP Precedence value. After you select this option, use the IP Precedence Value field to select the IP Precedence value to mark in packets that match the policy-class.
Mirror Interface	Select this option to copy the traffic stream to a specified egress port (physical or LAG) without bypassing normal packet forwarding. This action can occur in addition to any marking or policing action. It may also be specified along with a QoS queue assignment. Use the Interface menu to select the interface to which traffic is mirrored.

Table 317: Diffserv Policy Configuration Fields (continued)

Field	Description
Police Simple	<p>Select this option to enable the simple traffic policing style for the policy-class. The simple form of the police attribute uses a single data rate and burst size, resulting in two outcomes (conform and violate). After you select this option, configure the following policing criteria:</p> <ul style="list-style-type: none"> • Color Mode: The type of color policing used in DiffServ traffic conditioning. • Color Conform Class: For color-aware policing, packets in this class are metered against both the committed information rate (CIR) and the peak information rate (PIR). The class definition used for policing color awareness is only allowed to contain a single, non-excluded class match condition identifying one of the supported comparison fields: CoS, IP DSCP, IP Precedence, or Secondary COS. • Committed Rate (Kbps): The maximum allowed arrival rate of incoming packets for this class. • Committed Burst Size (Kbytes): The amount of conforming traffic allowed in a burst. • Conform Action: The action taken on packets that are considered conforming (below the police rate). • Violate Action: The action taken on packets that are considered non-conforming (above the police rate).
Police Single Rate	<p>Select this option to enable the single-rate traffic policing style for the policy-class. The single-rate form of the police attribute uses a single data rate and two burst sizes, resulting in three outcomes (conform, exceed, and violate). After you select this option, configure the following policing criteria:</p> <ul style="list-style-type: none"> • Color Mode: The type of color policing used in DiffServ traffic conditioning. • Color Conform Class: For color-aware policing, packets are metered against the committed information rate (CIR) and the peak information rate (PIR). The class definition used for policing color awareness is only allowed to contain a single, non-excluded class match condition identifying one of the supported comparison fields: CoS, IP DSCP, IP Precedence, or Secondary COS. This field is available only if one or more classes that meets the color-awareness criteria exist. • Color Exceed Class: For color-aware policing, packets are metered against the PIR only. • Committed Rate (Kbps): The maximum allowed arrival rate of incoming packets for this class. • Committed Burst Size (Kbytes): The amount of conforming traffic allowed in a burst. • Excess Burst Size (Kbytes): The amount of conforming traffic allowed to accumulate beyond the Committed Burst Size (Kbytes) value during longer-than-normal idle times. This value allows for occasional bursting. • Conform Action: The action taken on packets that are considered conforming (below the police rate). • Exceed Action: The action taken on packets that are considered to exceed the committed burst size but are within the excessive burst size. • Violate Action: The action taken on packets that are considered non-conforming (above the police rate).

Table 317: Diffserv Policy Configuration Fields (continued)

Field	Description
Police Two Rate	<p>Select this option to enable the two-rate traffic policing style for the policy-class. The two-rate form of the police attribute uses two data rates and two burst sizes. Only the smaller of the two data rates is intended to be guaranteed. After you select this option, configure the following policing criteria:</p> <ul style="list-style-type: none"> • Color Mode: The type of color policing used in DiffServ traffic conditioning. • Color Conform Class: For color-aware policing, packets are metered against the committed information rate (CIR) and the peak information rate (PIR). The class definition used for policing color awareness is only allowed to contain a single, non-excluded class match condition identifying one of the supported comparison fields: CoS, IP DSCP, IP Precedence, or Secondary COS. This field is available only if one or more classes that meets the color-awareness criteria exist. • Color Exceed Class: For color-aware policing, packets are metered against the PIR. • Committed Rate (Kbps): The maximum allowed arrival rate of incoming packets for this class. • Committed Burst Size (Kbytes): The amount of conforming traffic allowed in a burst. • Peak Rate (Kbps): The maximum peak information rate for the arrival of incoming packets for this class. • Excess Burst Size (Kbytes): The maximum size of the packet burst that can be accepted to maintain the Peak Rate (Kbps). • Conform Action: The action taken on packets that are considered conforming (below the police rate). • Exceed Action: The action taken on packets that are considered to exceed the committed burst size but are within the excessive burst size. • Violate Action: The action taken on packets that are considered non-conforming (above the police rate).
Redirect Interface	<p>Select this option to force a classified traffic stream to the specified egress port (physical port or LAG). Use the Interface field to select the interface to which traffic is redirected.</p>

Click **Refresh** to update the page with the most current data from the switch.

DiffServ Service Summary

Use the **DiffServ Service Summary** page to add DiffServ policies to interfaces, remove policies from interfaces, and edit policy-interface mappings.

To access this page, click **QoS > Diffserv > Service Summary** in the navigation menu.

Use the buttons to perform the following tasks:

- To add a policy to an interface, click **Add**.
- To edit a configured interface-policy association, select the entry to modify and click **Edit**.
- To remove one or more configured interface-policy associations, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.

Table 318: DiffServ Service Summary Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. Only interfaces that have an associated policy are listed in the table.
Direction	The traffic flow direction to which the policy is applied: <ul style="list-style-type: none"> • Inbound: The policy is applied to traffic as it enters the interface. • Outbound: The policy is applied to traffic as it exits the interface.
Status	The status of the policy on the interface. A policy is Up if DiffServ is globally enabled, and if the interface is administratively enabled and has a link. Otherwise, the status is Down.
Policy	The DiffServ policy associated with the interface.

When you click **Add** or **Edit**, the **Configure Service** window opens and allows you to configure DiffServ interface policies. Specifying 'None' for a policy has no effect when adding or editing interface policies. To remove an interface policy mapping, use the Remove button on the parent page. The following information describes the fields in this window.

Table 319: Configure Service Fields

Field	Description
Interface	Select an interface to associate with a policy.
Policy In	The menu lists all policies configured with a type of In. Select the policy to apply to traffic as it enters the interface.
Policy Out	The menu lists all policies configured with a type of Out. Select the policy to apply to traffic as it exits the interface.

Click **Refresh** to update the page with the most current data from the switch.

Diffserv Service Statistics

The **Diffserv Service Statistics** page displays service-level statistical information for all interfaces in the system to which a DiffServ policy has been attached.

To access this page, click **QoS > Diffserv > Service Statistics** in the navigation menu.

Table 320: Diffserv Service Statistics Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. The table displays all interfaces that have a DiffServ policy currently attached in a traffic flow direction.
Direction	The traffic flow direction to which the policy is applied: <ul style="list-style-type: none"> • In: The policy is applied to traffic as it enters the interface. • Out: The policy is applied to traffic as it exits the interface.
Status	The operational status of this service interface, either Up or Down.

Table 320: Diffserv Service Statistics Fields (continued)

Field	Description
Octets Offered	The total number of octets offered to all class instances in this service policy before their defined DiffServ treatment is applied. This is the overall count per-interface, per-direction.
Octets Discarded	The total number of octets discarded for all class instances in this service policy for any reason due to DiffServ treatment. This is the overall count per-interface, per-direction.
Octets Sent	The total number of octets forwarded for all class instances in this service policy after their defined DiffServ treatments were applied. In this case, forwarding means the traffic stream was passed to the next functional element in the data path, such as the switching or routing function of an outbound link transmission element. This is the overall count per-interface, per-direction.

Click **Refresh** to update the page with the most current data from the switch.

Diffserv Service Policy Statistics

The **Diffserv Service Policy Statistics** page displays class-oriented statistical information for the policy, which is specified by the interface and direction.

To access this page, click **QoS > Diffserv > Policy Statistics** in the navigation menu.

Table 321: Diffserv Service Policy Statistics Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. The table displays all interfaces that have a DiffServ policy currently attached in a traffic flow direction.
Direction	The traffic flow direction to which the policy is applied: <ul style="list-style-type: none"> • In: The policy is applied to traffic as it enters the interface. • Out: The policy is applied to traffic as it exits the interface.
Policy	The name of the policy currently attached to the interface.
Status	The operational status of the policy currently attached to the interface.
Class	The DiffServ class currently defined for the attached policy.
Octets Offered	The total number of octets offered to all class instances in this service policy before their defined DiffServ treatment is applied. This is the overall count per-interface, per-direction.
Octets Discarded	The total number of octets discarded for all class instances in this service policy for any reason due to DiffServ treatment. This is the overall count per-interface, per-direction.
Packets Offered	The total number of packets offered to all class instances in this service policy before their defined DiffServ treatment is applied. This is the overall count per-interface, per-direction.
Packets Discarded	The total number of packets discarded for all class instances in this service policy for any reason due to DiffServ treatment. This is the overall count per-interface, per-direction.

Click **Refresh** to update the page with the most current data from the switch.

A Configuration Examples

Configuring VLANs
Configuring Multiple Spanning Tree Protocol
Configuring VLAN Routing
Configuring 802.1X Network Access Control
Configuring Authentication Tiering
Configuring Differentiated Services for VoIP
IGMP and MLD Snooping Switches
Configuring Port Mirroring
Bidirectional Forwarding Detection

This appendix contains examples of how to configure selected features available in the 200 Series software. Each example contains procedures on how to configure the feature by using the web interface, and/or CLI, and/or *SNMP (Simple Network Management Protocol)*.



Note

Each configuration example starts from a factory-default configuration unless otherwise noted.

Configuring VLANs

Figure 6 shows a switch with four ports configured to handle the traffic for two VLANs. Port 1/0/2 handles traffic for both VLANs, while port 1/0/1 is a member of VLAN 2 only, and ports 1/0/3 and 1/0/4 are members of VLAN 3 only.

The following examples show how to create VLANs, assign ports to the VLANs, and assign a VLAN as the default VLAN to a port.

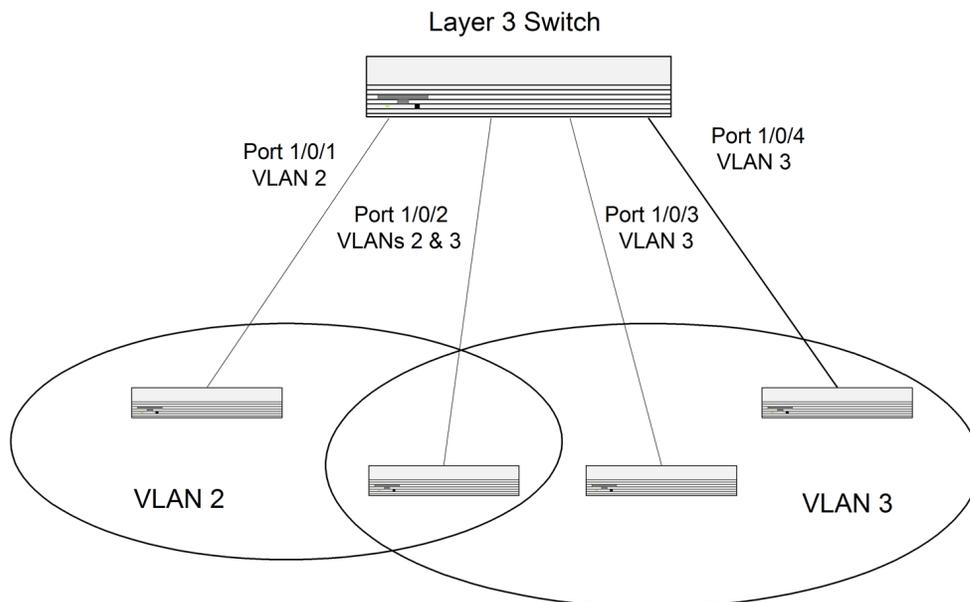


Figure 6: VLAN Example Network Diagram

Using the Web Interface to Configure VLANs

- 1 Access the **Switching > VLAN > Status** page.
- 2 Click **Add** to create a new VLAN.
- 3 Type 2-3 in the VLAN ID-Individual/Range field.
- 4 Click **Submit**.
- 5 From the **Port Configuration Page**, select VLAN 2 from the VLAN ID List.
- 6 From the Participation column in the interface table, select Include for ports 1/0/1 and 1/0/2 to specify that these ports are members of VLAN 2.
- 7 Select the interface checkbox and click **Edit**. Select the Tagging All checkbox to specify that frames will always be transmitted tagged from ports that are members of VLAN 2.
- 8 Click **Submit**.
- 9 Select VLAN 3 from the VLAN ID and Name List.
- 10 Select the Participate option in the VLAN field.
- 11 For ports 1/0/2, 1/0/3 and 1/0/4, select Include from the Participation menu to specify that these ports are members of VLAN 3.
- 12 Click **Submit**.
- 13 Go to the **Switching > VLAN > Port Configuration** page.
- 14 From the **Interface** menu, select 1/0/1.
- 15 In the Acceptable Frame Types field, select AdmitTaggedOnly to specify that untagged frames will be rejected on receipt.
- 16 Click **Submit**.
- 17 From the Interface menu, select 1/0/2.
- 18 In the Port VLAN ID field, enter 3 to assign VLAN 3 as the default VLAN for the port.

- 19 In the Acceptable Frame Types field, select AdmitTaggedOnly to specify that untagged frames will be rejected on receipt.
- 20 Click **Submit**.

Using the CLI to Configure VLANs

- 1 Create VLAN 2 and VLAN 3.

```
(Extreme 220) (Routing) #vlan database
vlan 2
vlan 3
exit
```

- 2 Assign ports 1/0/1 and 1/0/2 to VLAN2 and specify that untagged frames will be rejected on receipt.

```
(Extreme 220) (Routing) #Config
interface1/0/1
vlan participation include 2
vlan acceptframe vlanonly
exit
interface1/0/2
vlan participation include 2
vlan acceptframe vlanonly
```

- 3 While in interface config mode for port 1/0/2, assign VLAN3 as the default VLAN.

```
(Extreme 220) (Routing) (Interface 1/0/2)#vlan pvid 3
exit
```

- 4 Specify that frames will always be transmitted tagged from ports that are members of VLAN 2.

```
(Extreme 220) (Routing) (Config)#vlan port tagging all 2
exit
```

- 5 Assign the ports that will belong to VLAN 3.

Port 1/0/2 belongs to both VLANs, and port 1/0/1 can never belong to VLAN 3.

```
(Extreme 220) (Routing) #Config
interface1/0/2
vlan participation include 3
exit
interface1/0/3
vlan participation include 3
exit
interface1/0/4
vlan participation include 3
exit
exit
```

- 6 Specify that untagged frames will be accepted on port 1/0/4.

```
(Extreme 220) (Routing) #Config
interface1/0/4
vlan acceptframe all
exit
exit
```

Using the SNMP to Configure VLANs

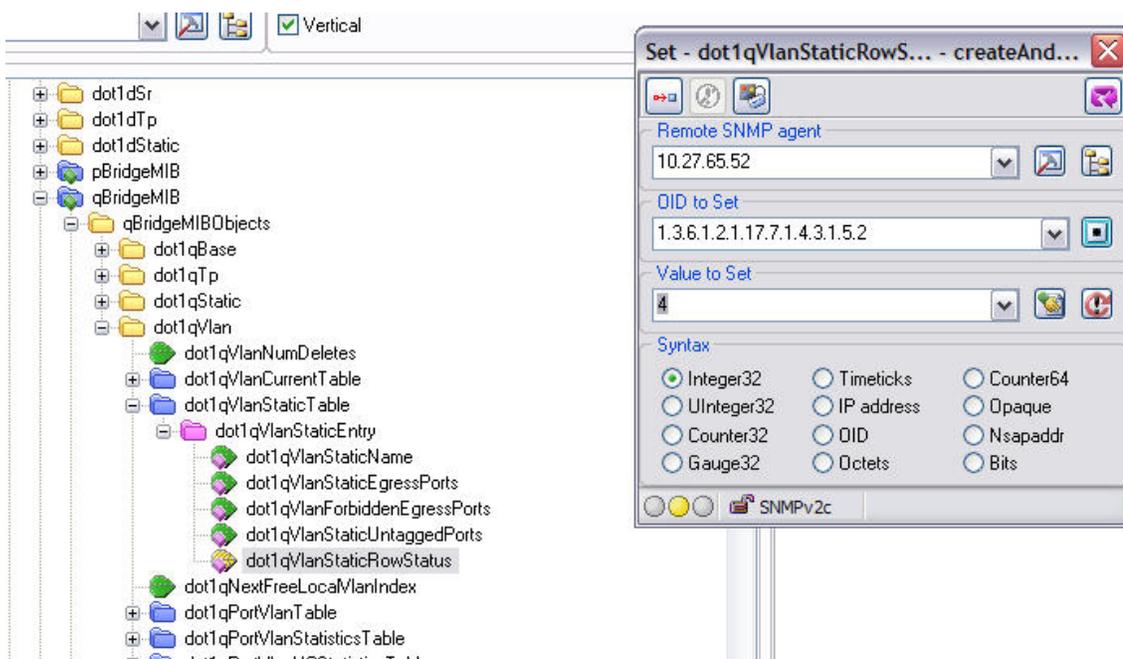
- 1 Use the objects in dot1qVlanStaticTable (in dot1qVlan in the QBRIDGE-MIB module) to create VLANs 2 and 3.

Set the dot1qVlanStaticRowStatus object to 'CreateandGo (4)' to create a VLAN. If the other parameters are not specified, simply specifying the dot1qVlanIndex and dot1qVlanStaticRowStatus is sufficient to create the VLAN.

The full path to the object is

iso(1).org(3).dod(6).internet(1).mgmt(2).mib-2(1).dot1dBridge(17).qBridgeMIB(7).

qBridgeMIBObjects(1).dot1qVlan(4).dot1qVlanStaticTable(3).dot1qVlanStaticEntry(1).dot1qVlanStaticRowStatus(5).



- 2 To assign ports 1/0/1 and 1/0/2 to VLAN2, retrieve the current dot1qStaticEgressPorts mask and append interfaces 1/0/1 and 1/0/2 to this mask by setting the first octet to 0xC0.

The dot1qVlanStaticEgressPorts bit mask can be constructed according to the following rules:

- Each octet within this value specifies a set of eight ports, with the first octet specifying ports (1-8), the second octet specifying ports (9-16), and so on.
- Within each octet, the most significant bit represents the lowest numbered port, and the least significant bit represents the highest numbered port. Thus, each port of the bridge is represented by a single bit within the value of this object. If that bit has a value of 1, then that port is included in the set of ports. The port is not included if its bit has a value of 0.

For example, if the switch has 12 ports and we want to add ports 1 and 4 in the VLAN and exclude all other ports, then the bit mask in hex will be 0x50 0x00.

- 3 To specify that frames will always be transmitted tagged from ports that are members of VLAN 2, use the dot1qVlanStaticUntaggedPorts object and set the value of the appropriate number of octets to 0.

Each octet represents eight ports, so for a 48-port switch, the first six octets would be zero.

- 4 To specify that ports 1/0/1 and 1/0/2 will only accept tagged frames and will reject untagged frames on receipt, set the dot1qPortAcceptableFrameTypes object to admitOnlyVlanTagged(2).
The object is in dot1qPortVlanEntry in the dot1qPortVlanTable.
- 5 To assign VLAN3 as the default VLAN for interface 1/0/2., set the value of dot1qPvid for 1/0/2 (instance 2) to 3.
- 6 To assign ports 1/0/2, 1/0/3, and 1/0/4 to VLAN3, retrieve the current dot1qStaticEgressPorts mask and append the interfaces to this mask by setting the first octet to 0x70.

Configuring Multiple Spanning Tree Protocol

This example shows how to enable IEEE 802.1s *MSTP (Multiple Spanning Tree Protocol)* on the switch and all of the ports and to set the bridge priority.

To make multiple switches be part of the same MSFP region, make sure the Force Protocol Version setting for all switches is IEEE 802.1s. Also, make sure the configuration name, digest key, and revision level are the same for all switches in the region.

Note



The digest key is generated based on the association of VLANs to different instances. To ensure the digest key is same, the mapping of VLAN to instance must be the same on each switch in the region. For example, if VLAN 10 is associated with instance 10 on one switch, you must associate VLAN 10 and instance 10 on the other switches.

Using the Web UI to Configure MSTP

- 1 Create VLANs 10 and 20.
 - a Access the **Switching > VLAN > Status** page.
 - b Click **Add** to create a VLAN.
 - c Select the VLAN ID-Individual option and enter 10.
 - d Click **Submit**.
 - e Repeat the steps to add VLAN 20.
- 2 Enable MSTP (IEEE 802.1s) on the switch and change the configuration name.
Changing the configuration name allows all the bridges that want to be part of the same region to join.
 - a Go to the **Switching > Spanning Tree > Switch** page.
 - b From the Spanning Tree Admin Mode menu, select **Enable**.
 - c In the **Configuration Name** field, enter `extreme`.
 - d Click **Submit**.

- 3 Create two MST instances.
 - a Go to the **Switching > Spanning Tree > MST** page.
 - b From the MST page, click **Add**.
 - c In the MST ID field, enter 10.
 - d Associate MST ID 10 with VLAN 10 and assign a bridge priority of 16384.
 - e Click **Submit**.
 - f Repeat the steps to create an MST instance with an ID of 20.
- 4 Use similar procedures to associate MST instance 20 to VLAN 20 and assign it a bridge priority value of 61440.

By using a lower priority for MST 20, MST 10 becomes the root bridge.
- 5 Force port 1/0/2 to be the root port for MST 20, which is the non-root bridge.
 - a Go to the **Switching > Spanning Tree > MST** page.
 - b From the MST ID menu, select 20.
 - c From the Interface menu, select 1/0/2.
 - d In the Port Priority field, enter 64.
 - e Click **Submit**.

Using the CLI to Configure MSTP

- 1 Create VLAN 10 and VLAN 20.


```
(Extreme 220) (Routing) #vlan database
vlan 10
vlan 20
exit
```
- 2 Enable spanning tree Globally


```
(Extreme 220) (Routing) #config
spanning-tree
```
- 3 Create MST instances 10 and 20.


```
spanning-tree mst instance 10
spanning-tree mst instance 20
```
- 4 Associate MST instance 10 to VLAN 10 and MST instance 20 to VLAN 20


```
spanning-tree mst vlan 10 10
spanning tree mst vlan 20 20
```
- 5 Change the name so that all the bridges that want to be part of the same region can form the region.


```
spanning-tree configuration name broadcom
```
- 6 Make the MST ID 10 bridge the root bridge by lowering the priority.


```
spanning-tree mst priority 10 16384
```
- 7 Change the priority of MST ID 20 to ensure the other bridge is the root bridge.


```
spanning-tree mst priority 20 61440
```
- 8 Enable STP on interface 1/0/1


```
interface 1/0/1
spanning-tree port mode
exit
```
- 9 Enable STP on interface 1/0/2


```
interface 1/0/2
spanning-tree port mode
```

- 10 On the non-root bridge change the priority to force port 1/0/2 to be the root port.

```
spanning-tree mst 20 port-priority 64
exit
```

Using SNMP to Configure MSTP

For instance 2, set the value to 64.

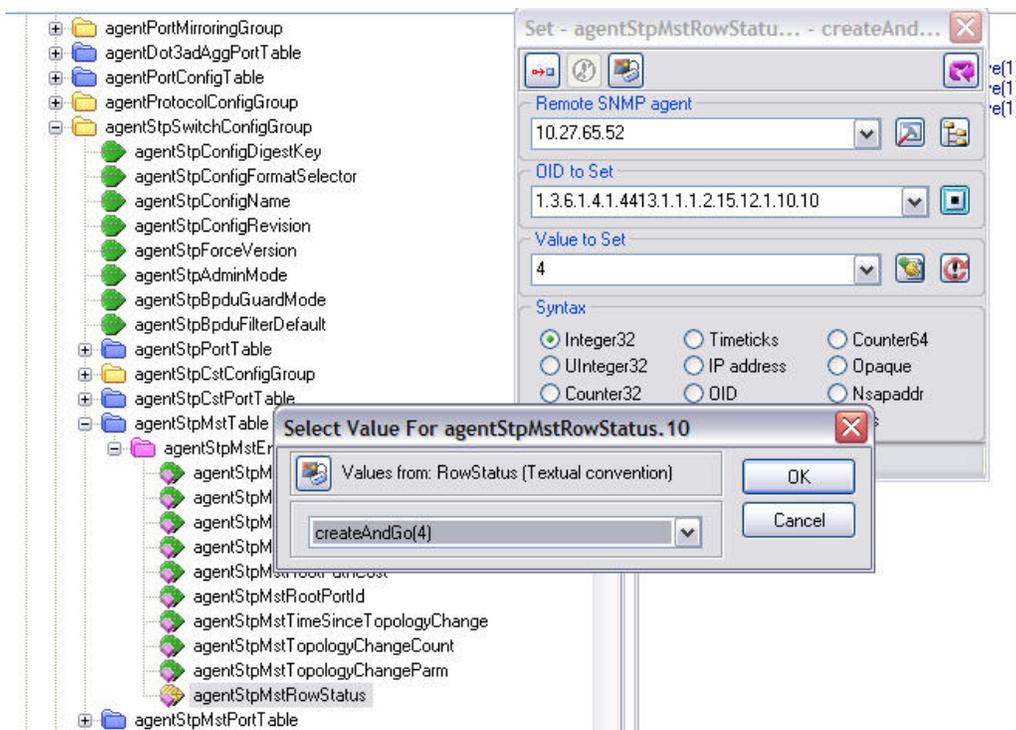
- 1 Use the objects in dot1qVlanStaticTable (in dot1qVlan in the QBRIDGE-MIB module) to create VLANs 10 and 20.
- 2 To enable spanning tree globally, set the agentStpAdminMode object in the 200 Series-SWITCHING-MIB module to enable (2).

The full path to the object is

iso(1).org(3).dod(6).internet(1).private(4).enterprises(1).broadcom(4413).

broadcomProducts(1).fastPath(1).fastPathSwitching(1).agentConfigGroup(2).agentStpSwitchConfigGroup(15).agentStpAdminMode(6).

- 3 Use the agentStpConfigName object in the agentStpSwitchConfigGroup to change the name so that all the bridges that want to be part of the same region can form the region.
- 4 Use the agentStpMstRowStatus object in the agentStpMstTable to create MST instances 10 and 20.



- 5 Use the agentStpMstBridgePriority object to set the bridge priorities for MST 10 and MST 20:
 - For MST ID 10, set the value to 16384 to make it the root bridge.
 - For MST ID 20, set the value to 61440 to ensure the other bridge is the root bridge.

- 6 Use the agentStpMstVlanRowStatusAssociate object in the agentStpMstVlanTable to associate MST instance 10 to VLAN 10 and MST instance 20 to VLAN 20.
 - For MST ID 20, the OID to set is 1.3.6.1.4.1.4413.1.1.1.2.15.14.1.1.10.10 (the final .10 is the VLAN ID)
 - For MST ID 20, the OID to set is 1.3.6.1.4.1.4413.1.1.1.2.15.14.1.1.20.20

Set the value to CreateAndGo (4).

- 7 Use the agentStpPortState in agentStpPortTable under agentStpSwitchConfigGroup to enable STP on interface 1/0/1 and interface 1/0/2.

For instance 1 and 2, set the value to enable (1).

- 8 Use the agentStpMstPortPriority object in agentStpMstPortTable to change the port priority on interface 1/0/2 to force the port to be the root port on the non-root bridge.

Configuring VLAN Routing

This section provides an example of how to configure 200 Series software to support VLAN routing. The configuration of the VLAN router port is similar to that of a physical port. The main difference is that, after the VLAN has been created, you must use the `show ip vlan` command to determine the VLAN's interface ID so that you can use it in the router configuration commands.

Figure 7 shows a Layer 3 switch configured for port routing. It connects two VLANs, with two ports participating in one VLAN, and one port in the other. The script shows the commands you would use to configure 200 Series software to provide the VLAN routing support shown in the diagram.

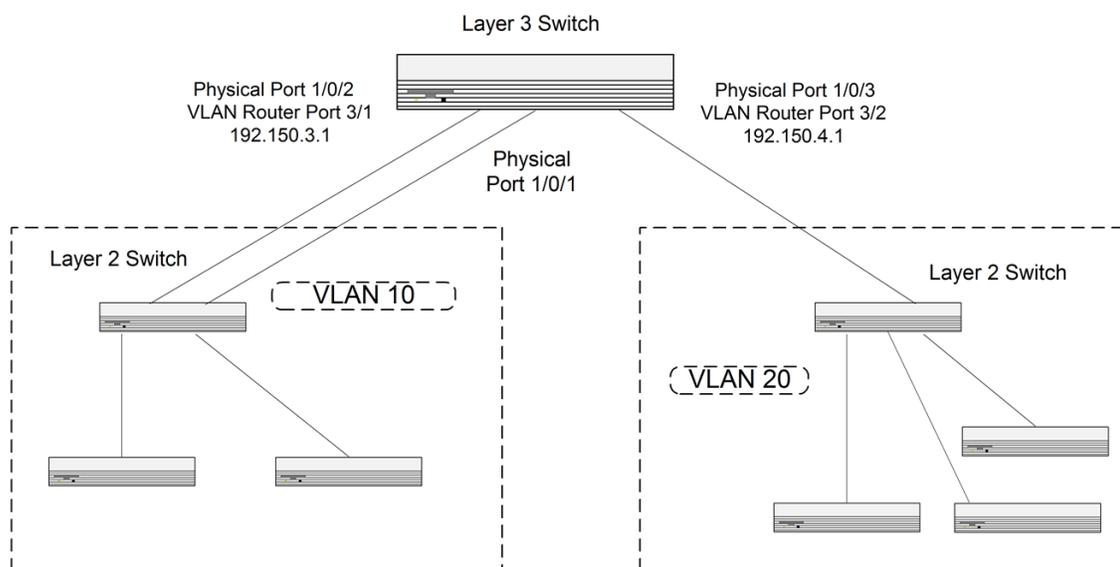


Figure 7: VLAN Routing Example Network Diagram

Using the CLI to Configure VLAN Routing

```
1 (Extreme 220) (Routing) #vlan database
   vlan 10
   vlan 20
   exit
```

- Configure ports 1/0/1, 1/0/2 as members of VLAN 10 and specify that untagged frames received on these ports will be assigned to VLAN 10.

```
config
interface 1/0/1
vlan participation include 10
vlan pvid 10
exit
interface 1/0/2
                vlan participation include 10
vlan pvid 10
exit
```

- Configure port 1/0/3 as a member of VLAN 20 and specify that untagged frames received on these ports will be assigned to VLAN 20

```
interface 1/0/3
                vlan participation include 20
vlan pvid 20
exit
exit
```

- Specify that all frames transmitted for VLANs 10 and 20 will be tagged.

```
config
vlan port tagging all 10
vlan port tagging all 20
exit
```

- Enable routing for the VLANs:

```
(Extreme 220) (Routing) #vlan database
vlan routing 10
vlan routing 20
exit
```

- View the logical interface IDs assigned to the VLAN routing interfaces.

```
(Extreme 220) (Routing) #show ip vlan
MAC Address used by Routing VLANs: 00:00:AA:12:65:12
                Logical
VLAN ID   Interface      IP Address      Subnet Mask
-----
10        0/4/1          0.0.0.0         0.0.0.0
20        0/4/2          0.0.0.0         0.0.0.0
```

As the output shows, VLAN 10 is assigned ID 0/4/1 and VLAN 20 is assigned ID 0/4/2

- Enable routing for the switch:

```
config
ip routing
exit
```

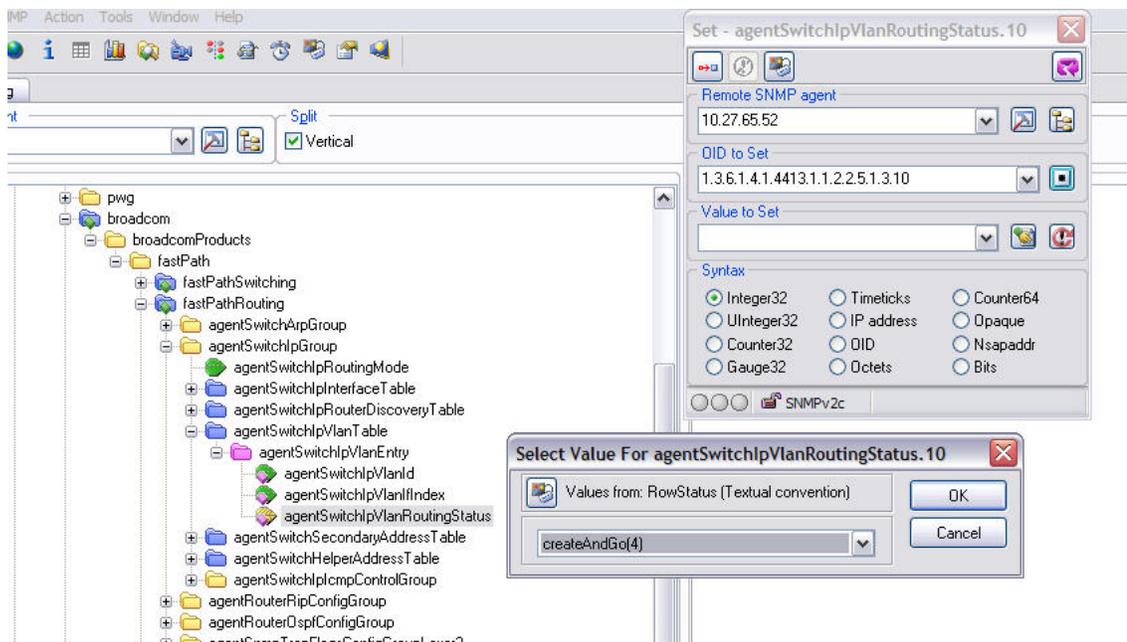
- Configure the IP addresses and subnet masks for the virtual router ports.

```
config
interface 0/4/1
ip address 192.150.3.1 255.255.255.0
exit
interface 0/4/2
                ip address 192.150.4.1 255.255.255.0
exit
exit
```

Using SNMP to Configure VLAN Routing

While setting the IP address for the VLAN interface, the `agentSwitchIpInterfaceIpAddress` and `agentSwitchIpInterfaceNetMask` should be set together.

- VLAN index 482 (VLAN 10): 192.150.3.1 255.255.255.0
 - VLAN index 483 (VLAN 20): 192.150.4.1 255.255.255.0
- 1 Use the dot1qVlanStaticRowStatus object in the dot1qVlanStaticTable to create VLAN 10 and VLAN 20.
 - 2 To configure VLAN membership, retrieve the current dot1qStaticEgressPorts mask and append the desired interfaces to the mask.
 - VLAN 10: 1/0/1 and 1/0/2
 - VLAN 20: 1/0/3
 - 3 To assign the PVID for an interface, use the dot1qPvid object.
 - 1/0/1: PVID 10
 - 1/0/2: PVID 10
 - 1/0/3: PVID 20
 - 4 To specify that all frames transmitted for VLANs 10 and 20 will be tagged, use the dot1qVlanStaticUntaggedPorts object and set the value of the appropriate number of octets to 0. Each octet represents eight ports, so for a 48-port switch, the first six octets would be zero.
 - 5 To enable routing for the VLANs, use the agentSwitchIpVlanRoutingStatus object in the agentSwitchIpVlanTable under agentSwitchIpGroup in fastPathRouting to set the value for VLAN 10 and VLAN 20 to CreateAndGo (4).



- 6 Walk the agentSwitchIpVlanIfIndex object to view the logical interface IDs assigned to the VLAN routing interfaces.
- 7 Set the agentSwitchIpRoutingMode object to enable (1) to enable routing for the switch:
- 8 Use the agentSwitchIpInterfaceIpAddress and agentSwitchIpInterfaceIpMask objects in the agentSwitchIpInterfaceTable to configure the IP addresses and subnet mask for the virtual router ports.

Configuring 802.1X Network Access Control

This example configures a single *RADIUS (Remote Authentication Dial In User Service)* server used for authentication and accounting at 10.10.10.10. The shared secret is configured to be secret. The switch is configured to require that the 802.1X access method is through a RADIUS server. IEEE 802.1X port-based access control is enabled for the system, and interface 1/0/1 is configured to be in force-authorized mode because this is where the RADIUS server and protected network resources are located.

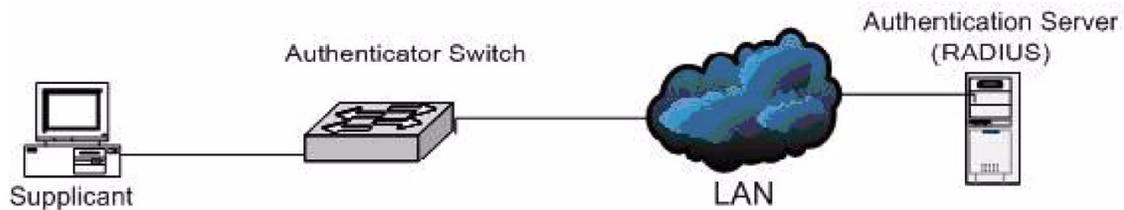


Figure 8: Switch with 802.1x Network Access Control

If a user, or supplicant, attempts to communicate via the switch on any interface except interface 1/0/1, the system challenges the supplicant for login credentials. The system encrypts the provided information and transmits it to the RADIUS server. If the RADIUS server grants access, the system sets the 802.1X port state of the interface to authorized, and the supplicant is able to access network resources.

Using the CLI to configure 802.1X Port-Based Access Control

- 1 Configure the *RADIUS* authentication server IP address.

```
(Extreme 220) (Config)#radius server host auth 10.10.10.10
```

- 2 Configure the RADIUS authentication server secret key.

```
(Extreme 220) (Config)#radius server key auth 10.10.10.10
```

You are prompted, and then re-prompted, to enter the secret key.

- 3 Configure the RADIUS accounting server IP address.

```
(Extreme 220) (Config)#radius server host acct 10.10.10.10
```

- 4 Configure the RADIUS accounting server secret key.

```
(Extreme 220) (Config)#radius server key acct 10.10.10.10
```

You are prompted, and then re-prompted, to enter the secret key.

- 5 Enable RADIUS accounting mode.

```
(Extreme 220) (Config)#radius accounting mode
```

- 6 Set IEEE 802.1X to use RADIUS as the AAA method.

```
(Extreme 220) (Config)#aaa authentication dot1x default radius
```

- 7 Enable 802.1X authentication on the switch.

```
(Extreme 220) (Config)#dot1x system-auth-control
```

- 8 Set the 802.1X mode for port 1/0/1 to Force Authorized.

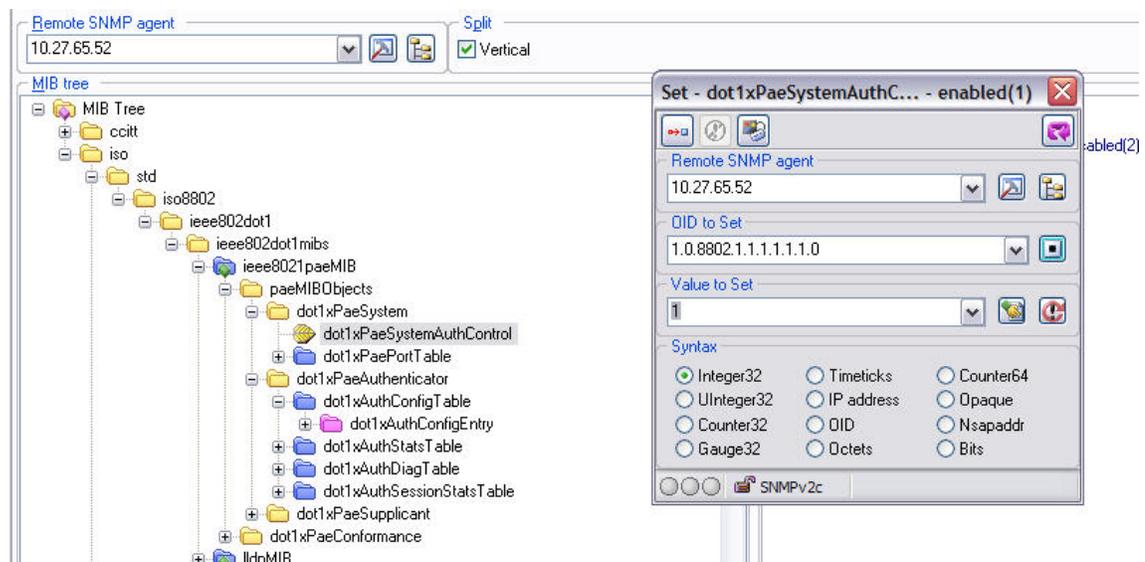
```
(Extreme 220) (Config)#interface 1/0/1
```

```
(Extreme 220) (Interface 1/0/1)#dot1x port-control force-authorized
```

```
(Extreme 220) (Interface 1/0/1)#exit
```

Using SNMP to configure 802.1X Port-Based Access Control

- 1 Use the agentRadiusServerStatus in the agentRadiusServerConfigTable under the FASTPATH-RADIUS-AUTH-CLIENT-MIB to create a new *RADIUS* server entry.
- 2 Use the agentRadiusServerAddress object to configure the RADIUS authentication server IP address as 10.10.10.10.
- 3 Use the agentRadiusServerSecret object to configure the RADIUS authentication server secret.
- 4 Use the agentRadiusAccountingStatus object in the agentRadiusAccountingConfigTable to create a RADIUS accounting server.
- 5 Use the agentRadiusAccountingServerAddress object to configure the RADIUS accounting server IP address. as 10.10.10.10.
- 6 Use the agentRadiusAccountingSecret object to configure the RADIUS accounting server secret.
- 7 Use the agentRadiusAccountingStatus object to enable RADIUS accounting mode.
- 8 Use the agentUserConfigDefaultAuthenticationList object in agentAuthenticationGroup in the FASTPATH-SWITCHING module to set RADIUS as the default login list for dot1x.
- 9 To enable 802.1X authentication on the switch, set the dot1xPaeSystemAuthControl object in the IEEE8021-PAE-MIB module to enable (1).



- 10 To set the 802.1X mode for port 1/0/1 to Force Authorized, use the agentDot1xPortControlMode object in the agentDot1xPortConfigTable, which is in FASTPATH-DOT1X-ADVANCED-FEATURES-MIB.

Configuring Authentication Tiering

Authentication Tiering can be configured in either of the following ways:

- [Configuring Authentication Tiering Using the Web Interface](#) on page 342
- [Configuring Authentication Tiering Using the CLI](#) on page 343

Configuring Authentication Tiering Using the Web Interface

To configure Authentication Tiering through the web interface:

- 1 Access the **Security > Authentication Manager > Configuration** page.
The **Authentication Manager Configuration** window opens.
- 2 Select **Enable**.
- 3 Click **Submit**.
- 4 Access the **Security > Authentication Manager > Authentication Tiering** page.
The **Authentication Tiering** window opens.
- 5 Select interface 1/0/3 checkbox and click **Edit**.
The **Edit Authentication Tiering** window opens.
- 6 Type **10000** in the Re-Authentication Timer field.
- 7 In the Configured Method Order box, move **Dot1x**, **MAB**, and **Captive Portal** to the Order box by selecting the method and clicking the **>** button.

**Note**

Captive portal is not supported in this version of the product.

- 8 In the Configured Method Priority box, move **Dot1x** and **Captive Portal** to the Priority box by selecting the method and clicking the **>** button.
- 9 Click **Submit**.

Configuring Authentication Tiering Using the CLI

To configure Authentication Tiering using the CLI:

- 1 Enable Authentication Tiering globally.

```
config
authentication enable
exit
```

- 2 Configure the authentication order, priority, and restart timer on interface 1/0/3.

**Note**

Although `captive-portal` is a command option, captive portal is not supported in this version of the product.

```
config
interface 1/0/3
authentication order dot1x mab captive-portal
authentication priority captive-portal dot1x
authentication restart 10000
exit
exit
```

Configuring Differentiated Services for VoIP

One of the most valuable uses of DiffServ is to support Voice over IP (VoIP). VoIP traffic is inherently time-sensitive: for a network to provide acceptable service, a guaranteed transmission rate is vital. This example shows one way to provide the necessary quality of service: how to set up a class for UDP traffic, have that traffic marked on the inbound side, and then expedite the traffic on the outbound side. The configuration script is for Router 1 in [Figure 9](#). A similar script should be applied to Router 2.

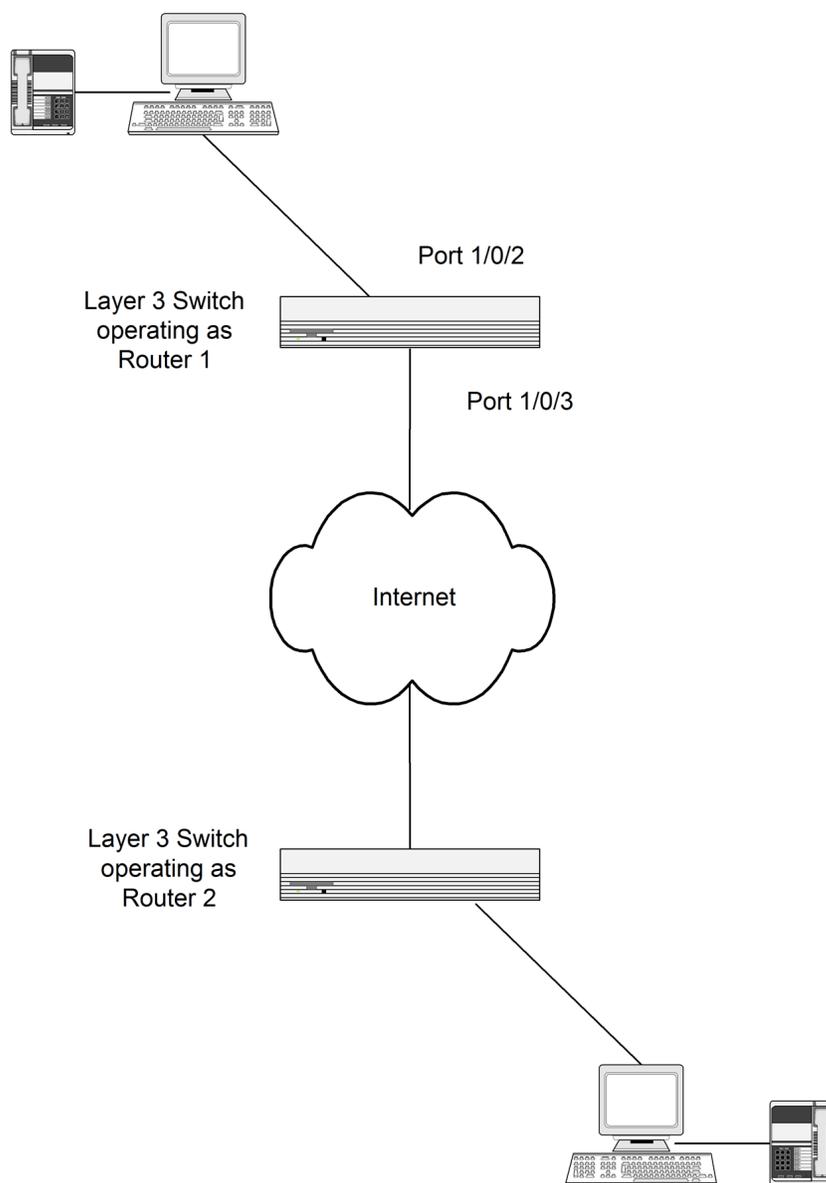


Figure 9: DiffServ VoIP Example Network Diagram

Using the CLI to Configure DiffServ VoIP Support

- 1 Enter Global Config mode. Set queue 5 on all ports to use strict priority mode. This queue shall be used for all VoIP packets. Activate DiffServ for the switch.

```
(Extreme 220) (Routing) #config
cos-queue strict 5
diffserv
```

- 2 Create a DiffServ classifier named 'class_voip' and define a single match criterion to detect UDP packets. The class type match-all indicates that all match criteria defined for the class must be satisfied in order for a packet to be considered a match.

```
class-map match-all class_voip
match protocol udp
exit
```

- 3 Create a second DiffServ classifier named 'class_ef' and define a single match criterion to detect a DiffServ code point (DSCP) of 'EF' (expedited forwarding). This handles incoming traffic that was previously marked as expedited elsewhere in the network.

```
class-map match-all class_ef
match ip dscp ef
exit
```

- 4 Create a DiffServ policy for inbound traffic named 'pol_voip', then add the previously created classes 'class_ef' and 'class_voip' as instances within this policy.

This policy handles incoming packets already marked with a DSCP value of 'EF' (per 'class_ef' definition), or marks UDP packets per the 'class_voip' definition) with a DSCP value of 'EF'. In each case, the matching packets are assigned internally to use queue 5 of the egress port to which they are forwarded.

```
policy-map pol_voip in
class class_ef
assign-queue 5
exit
class class_voip
mark ip-dscp ef
assign-queue 5
exit
exit
```

- 5 Attach the defined policy to an inbound service interface.

```
interface 1/0/2
service-policy in pol_voip
exit
exit
```

Using SNMP to Configure DiffServ VoIP Support

- 1 Use the agentDiffServGenStatusAdminMode object in agentDiffServGenStatusGroup under fastPathQOSDiffServPrivate in the FASTPATH-QOS-DIFFSERV-PRIVATE-MIB module to activate DiffServ for the switch.
- 2 To set queue 5 on all ports to use strict priority mode, use the agentCosQueueSchedulerType in the agentCosQueueTable in the FASTPATH-QOS-COS-MIB module. This queue is used for all VoIP packets.
- 3 Use the agentDiffServClassRowStatus object in the agentDiffServClassTable to create two new DiffServ instances. Set the value to CreateAndGo (4).
- 4 Use the agentDiffServClassName in the agentDiffServClassTable to name the first DiffServ classifier "class_voip" and the second classifier "class_ef."
- 5 Use the agentDiffServClassType in the agentDiffServClassTable to set the class type for each classifier to All (1).
- 6 Use the agentDiffServClassRuleMatchEntryType in the agentDiffServClassRuleTable to set class_voip to match a protocol (9) and class_ef to match an IP DSCP value (6).
- 7 For class_voip, define a single match criterion to detect UDP packets by setting the agentDiffServClassRuleMatchProtocolNum in the agentDiffServClassRuleTable to 17.
- 8 Use the agentDiffServClassRuleMatchIpDscp object in the agentDiffServClassRuleTable to define a single match criterion to detect a DSCP of EF (46). This handles incoming traffic that was previously marked as expedited elsewhere in the network.
- 9 Use the agentDiffServPolicyRowStatus object in the agentDiffServPolicyTable to create a DiffServ policy. Set the value to CreateAndGo (4).

- 10 Use the agentDiffServPolicyType object to set the policy direction so that it applies to inbound (I) traffic.
- 11 Use the agentDiffServPolicyName object to name the new DiffServ instance “pol_voip.”
- 12 Use the agentDiffServPolicyInstRowStatus object in the agentDiffServPolicyInstTable to create new instances that will be associated with the previously created classes (class_ef and class_voip).
- 13 Use the agentDiffServPolicyInstClassIndex object to associate class_ef and class_voip with the policy instances.
- 14 Use the agentDiffServPolicyAttrRowStatus object in the agentDiffServPolicyAttrTable to create three instances.
- 15 Use the agentDiffServPolicyAttrStmtAssignQueueId to set the queue value for instances 1.1.1 and 1.2.2 to 5, so that matching packets are assigned internally to use queue 5 of the egress port to which they are forwarded.
- 16 Use the agentDiffServPolicyAttrStmtMarkIpDscpVal object to set the value of instance 1.2.1 to 46, which marks UDP packets (per the class_voip definition) with a DSCP value of EF.
- 17 Create an instance for the interface that will have the policy attached by using the agentDiffServServiceRowStatus object in the agentDiffServServiceTable. For example, to create an instance for interface 1/0/2, set 2.1 to CreateAndGo (4).
- 18 Attach the policy to the interface instance by using the agentDiffServServicePolicyIndex object. Set the value of the instance to 1.

IGMP and MLD Snooping Switches

A snooping switch can be configured to receive *IGMP (Internet Group Management Protocol)* packets in a subnet and identify ports on which interested IP multicast listeners are present. It also identifies the ports on which multicast routers are attached, and these are the likely ports on which IP multicast sources are present. This section describes how the snooping switch handles IGMP messages, addresses considerations for IGMP packet and IP multicast traffic forwarding, and provides information about support for IGMP and MLD versions.

Query Processing

When a port receives an *IGMP* query, the snooping switch identifies the receiving port as a multicast router-attached port. The switch maintains the list of multicast router-attached ports on a per-VLAN basis. If the port does not receive periodic IGMP queries, the learned entries maintained in the list expire after a configured interval. The snooping switch stores the last received query on a VLAN because the switch uses this to calculate the max response time value during IGMP Leave message processing.

The snooping switch treats Distance Vector Multicast Routing Protocol (DVMRP) probe messages and Protocol Independent Multicast (PIM) versions 1 and 2 hello messages that it receives similar to IGMP queries by adding the interfaces that receive these messages to the list of multicast router-attached ports.

Group Registration

Multicast listeners can register to an IP multicast group by sending an *IGMP* Report message in response to a general query from a multicast router or by sending an unsolicited IGMP Report message. When the snooping switch processes an IGMP Report message, it creates an entry in the Layer 2

multicast forwarding table for the requested multicast group. Each entry contains a unique VLAN and multicast group combination along with a list of ports on which the IGMP Report was received. Multicast router-attached ports discovered during query processing on the incoming VLAN are automatically added to the newly created Layer 2 multicast forwarding entry.

The created entries expire if no additional IGMP Report messages are received for that multicast group, VLAN, and received port combination. The snooping switch administrator can configure the group expiry on a per-VLAN basis. If all the host registrations expire for a Layer 2 multicast forwarding entry, the entry is removed from the table.

Group Leave

Multicast listeners can opt to voluntarily leave a group by sending an *IGMP* Leave message or by not responding to the periodic IGMP queries sent by the multicast router. Upon receiving an IGMP Leave message, the snooping switch sends a group specific query on the received port to solicit IGMP Reports from other interested hosts on the same network segment. The snooping switch waits for the interval specified by the last received query packet (max response time) to receive a response for the Leave query. If there is no response, the port is removed from the Layer 2 multicast forwarding entry. If no querier information is available, a configured value is used. If an IGMP Report is received, the entry remains the same.

Alternatively, the administrator can configure the snooping switch to remove the interface that received the IGMP Leave message from the Layer 2 multicast forwarding entry immediately upon processing the message. No IGMP Leave query is sent in this scenario. Configuring the immediate leave is useful in situations where instantaneous control of group registrations is required, which results in better bandwidth control.

IGMP Packet Forwarding Considerations

The snooping switch forwards received *IGMP* Report and Leave messages only to multicast router-attached ports in that VLAN. IGMP queries are forwarded to all member interfaces of the VLAN.

The snooping switch is aware of link-layer changes caused by spanning tree operations. When a port is enabled or disabled by spanning tree, a general query is generated by the root bridge. This Topology Change Notification query is sent to all non-multicast router-attached ports of the root bridge, which aids in updating Layer 2 multicast forwarding entries faster so that network disruptions are felt only momentary.

The snooping switch processes all IGMP messages and drops invalid IGMP and MLD messages. Any unrecognized IGMP or MLD messages are forwarded in the VLAN. When the snooping switch originates an IGMP query (leave processing or TCN), it does not alter the version number or fields. The snooping switch leaves this information the same as the query information it received most recently on that VLAN.

IP Multicast Data Forwarding Considerations

When processing a packet whose destination MAC address is a multicast address, an IEEE standard bridge forwards a copy of the packet to each of the remaining network interfaces that are members of the same VLAN.

By default, unregistered multicast data packets are flooded to all ports in the VLAN.

By creating static Layer 2 multicast forwarding entries, multicast groups can be registered, and data can be forwarded only to selected ports.

Version Compatibility

The following table outlines the *IGMP*/MLD versions the 200 Series snooping switch supports.

Table 322: IGMP/MLD Version Support

Protocol Version	Support
IGMPv1	Yes
IGMPv2	Yes
IGMPv3	No
MLDv1	Yes
MLDv2	No

Snooping Switch Restrictions

This section describes the *IGMP* and MLD Snooping implementation on a 200 Series-based snooping switch.

MAC Address-Based Multicast Group

The L2 multicast forwarding table (built using IGMPV2/V1 reports) consists of the IP Multicast group MAC address. For IPv4 multicast groups, 16 IP multicast group addresses map to the same multicast MAC address. For example, 224.1.1.1 and 225.1.1.1 map to the MAC address 01:00:5E:01:01:01, and IP addresses in the range [224-239].3.3.3 map to 01:00:5E:03:03:03. As a result, if a host requests 225.1.1.1 using IGMPv2 or IGMPv1, then it might receive multicast traffic of group 226.1.1.1 as well.

IGMP Snooping in a Multicast Router

IGMP snooping is a Layer 2 feature and is achieved by using the Layer 2 multicast forwarding table. However when multicast routing is enabled on a 200 Series switch, Layer 2 multicast forwarding entries do not affect multicast data forwarding. Instead, corresponding IP multicast table entries need to be created to achieve similar behavior.

On a multicast router, for IGMP snooping to be functional, any multicast routing protocol needs to be operationally enabled on the routing interface. IGMP snooping also needs to be enabled on the VLAN corresponding to the routing interface. Note that IGMP snooping behavior will not be functional on VLANs that are not enabled for VLAN routing.

Configuring IGMP and MLD Snooping

The command-line interface (CLI) includes several commands that are used to configure the *IGMP* and MLD snooping features. For more information about each command, and for information about

commands that are not described in this section, see [ExtremeSwitching 200 Series: Command Reference Guide](#).

Enabling IGMP Snooping

To globally enable *IGMP* snooping on the switch enter Global Configuration mode and use the `set igmp` command, for example:

```
console(config) #set igmp
```

To enable IGMP snooping on an interface, enter Interface Configuration mode and use the `set igmp` command, for example:

```
console(config) #interface 1/0/1console(config-if-1/0/1) #set igmp
```

To enable IGMP snooping on a VLAN, enter VLAN Config mode and use the `set igmp vlan_id` command. The following example enables IGMP snooping on VLAN 10:

```
console #vlan databaseconsole(config-vlan) #ip igmp 10
```

Configuring IGMP Snooping Parameters

The following example shows how to configure the group membership interval on an interface (Interface Config mode):

```
console(Interface 1/0/1) #set igmp groupmembership-interval 250
```

The following example shows how to configure the group membership interval on VLAN 10 (VLAN Config mode):

```
console(Vlan) #set igmp groupmembership-interval 10 250
```

The following example shows how to configure the max response interval on an interface (Interface Config mode):

```
console(Interface 1/0/1) #set igmp maxresponse 10
```

The following example shows how to configure the max response interval on VLAN 10 (VLAN Config mode):

```
console(Vlan) #set igmp maxresponse 10 10
```

The following example shows how to enable fast leave mode on VLAN 10 (VLAN Config mode):

```
console(Vlan) #set igmp fast-leave 10
```

The following example shows how to configure the multicast router attached ports expiry interval on an interface (Interface Config mode):

```
console(Interface 1/0/1) #set igmp mcrtreptime 60
```

The following example shows how to configure the multicast router attached ports expiry interval on VLAN 10 (VLAN Config mode):

```
console(config-vlan) #set igmp mcrtreptime 10 60
```

Display IGMP Snooping Information

The following example shows how to display the *IGMP* snooping groups:

```
console#show mac-address-table igmpsnooping
```

Description	Interfaces	VLAN ID	MAC Address	Type
1	01:00:5E:01:02:03	Dynamic	Network Assist	Fwd: 1/0/2

Configuring Static Multicast Forwarding Entries

The following example shows how to create a static multicast forwarding entry for VLAN 1 and multicast MAC address 01:00:5E:11:22:33, associate it with the destination port 1/0/2 and the source port 1/0/4.

```
(Extreme 220) (Config)#macfilter 01:00:5e:11:22:33 1
(Extreme 220)#interface 1/0/2
(Extreme 220) (Interface 1/0/2)#macfilter adddest 01:00:5e:11:22:33 1
(Extreme 220) (Interface 1/0/2)#exit
(Extreme 220)#interface 1/0/4
(Extreme 220) (Interface 1/0/4)#macfilter addsrc 01:00:5e:11:22:33 1
(Extreme 220) (Interface 1/0/4)#exit
(Extreme 220)#show mac-address-table multicast
```

VLAN ID	MAC Address	Source	Type	Description	Interface	Fwd Interface
1	01:00:5E:11:22:33	Filter	Static	Mgmt Config	Fwd: 1/0/2	Fwd: 1/0/2

```
(Extreme 220)#show mac-address-table static all
```

MAC Address	VLAN ID	Source Port (s)	Destination Port (s)
01:00:5E:11:22:33	1	1/0/4	1/0/2

```
(Extreme 220)#show mac-address-table multicast 01:00:5e:11:22:33 1
```

VLAN ID	MAC Address	Source	Type	Description	Interface	Fwd Interface
1	01:00:5E:11:22:33	Filter	Static	Mgmt Config	Fwd: 1/0/2	Fwd: 1/0/2

Configuring Port Mirroring

Port mirroring is used to monitor the network traffic that a port sends and receives. The Port Mirroring feature creates a copy of the traffic that the source port handles and sends it to a destination port. The source port is the port that is being monitored. The destination port is monitoring the source port. The destination port is where you would connect a network protocol analyzer to learn more about the traffic that is handled by the source port.

A port monitoring session includes one or more source ports that mirror traffic to a single destination port. 200 Series software supports a single port monitoring session. LAGs (port channels) cannot be used as the source or destination ports.

For each source port, you can specify whether to mirror ingress traffic (traffic the port receives, or RX), egress traffic (traffic the port sends, or TX), or both ingress and egress traffic.

The packet that is copied to the destination port is in the same format as the original packet on the wire. This means that if the mirror is copying a received packet, the copied packet is VLAN tagged or

untagged as it was received on the source port. If the mirror is copying a transmitted packet, the copied packet is VLAN tagged or untagged as it is being transmitted on the source port.

After you configure the port mirroring session, you can enable or disable the administrative mode of the session to start or stop the probe port from receiving mirrored traffic.

Bidirectional Forwarding Detection



Note

Bidirectional Forwarding Detection (BFD) configuration can be performed only by using the CLI. web UI and *SNMP* configuration options are not supported for the BFD feature.

In a network device, BFD is presented as a service to its user applications, providing them options to create and destroy a session with a peer device and reporting upon the session status. On 200 Series switches, *BGP (Border Gateway Protocol)* can use BFD for monitoring of their neighbors' availability in the network and for fast detection of connection faults with them.

BFD uses a simple 'hello' mechanism that is similar to the neighbor detection components of some well-known protocols. It establishes an operational session between a pair of network devices to detect a two-way communication path between them and serves information regarding it to the user applications. The pair of devices transmits BFD packets between them periodically, and if one stops receiving peer packets within detection time limit it considers the bidirectional path to have failed. It then notifies the application protocol using its services.

BFD allows each device to estimate how quickly it can send and receive BFD packets to agree with its neighbor upon how fast detection of failure could be done.

BFD can operate between two devices on top of any underlying data protocol (network layer, link layer, tunnels, etc.) as payload of any encapsulating protocol appropriate for the transmission medium. The 200 Series implementation works with IPv4 and IPv6 networks and supports IPv4/v6 address-based encapsulations.

Configuring BFD

The following command sequence enables BFD and configures session parameters:

- 1 First, globally enable BFD:

```
(Router)#configure
(Router) (Config)# feature bfd
```

- 2 Configure session settings. These can be configured globally or on a per-interface basis.

```
(Router) (Config)#bfd interval 100 min_rx 200 multiplier 5
(Router) (Config)#bfd slow-timer 1000
```

- The argument **interval** refers to the desired minimum transmit interval, the minimum interval that the user wants to use while transmitting BFD control packets (in ms).
- The argument **min_rx** refers to the required minimum receive interval, the minimum interval at which the system can receive BFD control packets (in ms).

- The argument **multiplier** specifies the number of BFD control packets to be missed in a row to declare a session down.
 - The `slow-timer` command sets up the BFD required echo receive interval preference value (in ms). This value determines the interval the asynchronous sessions use for BFD control packets when the echo function is enabled. The slow-timer value is used as the new control packet interval, while the echo packets use the configured BFD intervals.
- 3 Configure BGP to use BFD for fast detection of faults between neighboring devices. A neighboring device IP address

```
(Router) (Config)#router bgp
(Router) (Config-router)# neighbor 172.16.11.6 fall-over bfd
(Router) (Config-router)# exit
```

Glossary

ACE

An Access Control Entry outlines the permissions associated with a particular ACE ID. There are three ACE frame types (Ethernet Type, ARP, and IPv4) and two ACE actions (permit and deny). The ACE also contains many parameter options that are available for individual application.

ACL

An Access Control List is a mechanism for filtering packets at the hardware level. Packets can be classified by characteristics such as the source or destination MAC, IP address, IP type, or QoS queue. Once classified, the packets can be forwarded, counted, queued, or dropped.

ad hoc mode

An 802.11 networking framework in which devices or stations communicate directly with each other, without the use of an AP.

ARP

Address Resolution Protocol is part of the TCP/IP suite used to dynamically associate a device's physical address (MAC address) with its logical address (IP address). The system broadcasts an ARP request, containing the IP address, and the device with that IP address sends back its MAC address so that traffic can be transmitted.

ATM

Asynchronous Transmission Mode is a start/stop transmission in which each character is preceded by a start signal and followed by one or more stop signals. A variable time interval can exist between characters. ATM is the preferred technology for the transfer of images.

BGP

Border Gateway Protocol is a router protocol in the IP suite designed to exchange network reachability information with BGP systems in other autonomous systems. You use a fully meshed configuration with BGP.

BGP provides routing updates that include a network number, a list of ASs that the routing information passed through, and a list of other path attributes. BGP works with cost metrics to choose the best available path; it sends updated router information only when one host has detected a change, and only the affected part of the routing table is sent.

BGP communicates within one AS using Interior BGP (IBGP) because BGP does not work well with IGP. Thus the routers inside the AS maintain two routing tables: one for the IGP and one for IBGP. BGP uses exterior BGP (EBGP) between different autonomous systems.

BPDU

In STP, a Bridge Protocol Data Unit is a packet that initiates communication between devices. BPDU packets contain information on ports, addresses, priorities, and costs and they ensure that the data ends up where it was intended to go. BPDU messages are exchanged across bridges to detect loops in a network topology. The loops are then removed by shutting down selected bridge interfaces and placing redundant switch ports in a backup, or blocked, state.

BSS

Basic Service Set is a wireless topology consisting of one access point connected to a wired network and a set of wireless devices. Also called an infrastructure network. See also [*IBSS \(Independent Basic Service Set\)*](#).

CDP

Cisco Discovery Protocol is a proprietary Data Link Layer protocol that shares information about other directly connected Cisco equipment, such as operating system versions and IP addresses.

Chalet

Chalet is a web-based user interface for setting up and viewing information about a switch, removing the need to enter common commands individually in the CLI.

CHAP

Challenge-Handshake Authentication Protocol is one of the two main authentication protocols used to verify a user's name and password for PPP Internet connections. CHAP is more secure because it performs a three-way handshake during the initial link establishment between the home and remote machines. It can also repeat the authentication anytime after the link has been established.

CLI

Command Line Interface. The CLI provides an environment to issue commands to monitor and manage switches and wireless appliances.

CoS

Class of Service specifies the service level for the classified traffic type.

Data Center Connect

DCC, formerly known as DCM (Data Center Manager), is a data center fabric management and automation tool that improves the efficiency of managing a large virtual and physical network. DCC provides an integrated view of the server, storage, and networking operations, removing the need to use multiple tools and management systems. DCC automates VM assignment, allocates appropriate network resources, and applies individual policies to various data objects in the switching fabric (reducing VM sprawl). Learn more about DCC at <http://www.extremenetworks.com/product/data-center-connect/>.

DHCP

Dynamic Host Configuration Protocol allows network administrators to centrally manage and automate the assignment of IP addresses on the corporate network. DHCP sends a new IP address when a computer is plugged into a different place in the network. The protocol supports static or dynamic IP addresses and can dynamically reconfigure networks in which there are more computers than there are available IP addresses.

DNS

A Domain Name Server is used to translate domain names to IP addresses. Although the Internet is based on IP addresses, names are easier to remember and work with. All these names must be translated back to the actual IP address and the DNS servers do so.

DoS attack

Denial of Service attacks occur when a critical network or computing resource is overwhelmed so that legitimate requests for service cannot succeed. In its simplest form, a DoS attack is indistinguishable

from normal heavy traffic. ExtremeXOS software has configurable parameters that allow you to defeat DoS attacks.

DSA

Digital Signature Algorithm is a Federal Information Processing Standard for digital signatures. It was proposed by the National Institute of Standards and Technology (NIST) in August 1991 for use in their Digital Signature Standard (DSS) and adopted as FIPS 186 in 1993. Four revisions to the initial specification have been released: FIPS 186-1 in 1996, FIPS 186-2 in 2000, FIPS 186-3 in 2009, and FIPS 186-4 in 2013.

DSSS

Direct-Sequence Spread Spectrum is a transmission technology used in Local Area Wireless Network (LAWN) transmissions where a data signal at the sending station is combined with a higher data rate bit sequence, or chipping code, that divides the user data according to a spreading ratio. The chipping code is a redundant bit pattern for each bit that is transmitted, which increases the signal's resistance to interference. If one or more bits in the pattern are damaged during transmission, the original data can be recovered due to the redundancy of the transmission. (Compare with [*FHSS \(Frequency-Hopping Spread Spectrum\)*](#).)

EAP-TLS/EAP-TTLS

EAP-TLS Extensible Authentication Protocol - Transport Layer Security. A general protocol for authentication that also supports multiple authentication methods, such as token cards, Kerberos, one-time passwords, certificates, public key authentication and smart cards.

IEEE 802.1x specifies how EAP should be encapsulated in LAN frames.

In wireless communications using EAP, a user requests connection to a WLAN through an access point, which then requests the identity of the user and transmits that identity to an authentication server such as RADIUS. The server asks the access point for proof of identity, which the access point gets from the user and then sends back to the server to complete the authentication.

EAP-TLS provides for certificate-based and mutual authentication of the client and the network. It relies on client-side and server-side certificates to perform authentication and can be used to dynamically generate user-based and session-based WEP keys.

EAP-TTLS (Tunneled Transport Layer Security) is an extension of EAP-TLS to provide certificate-based, mutual authentication of the client and network through an encrypted tunnel, as well as to generate dynamic, per-user, per-session WEP keys. Unlike EAP-TLS, EAP-TTLS requires only server-side certificates.

(See also [*PEAP \(Protected Extensible Authentication Protocol\)*](#).)

ECMP

Equal Cost Multi Paths is a routing algorithm that distributes network traffic across multiple high-bandwidth OSPF, BGP, IS-IS, and static routes to increase performance. The Extreme Networks implementation supports multiple equal cost paths between points and divides traffic evenly among the available paths.

ESRP

Extreme Standby Router Protocol is an Extreme Networks-proprietary protocol that provides redundant Layer 2 and routing services to users.

Extreme Application Analytics

EAA, formerly Purview™, is a network powered application analytics and optimization solution that captures and analyzes context-based application traffic to deliver meaningful intelligence about applications, users, locations, and devices. EAA provides data to show how applications are being used. This can be used to better understand customer behavior on the network, identify the level of user engagement, and assure business application delivery to optimize the user experience. The software also provides visibility into network and application performance allowing IT to pinpoint and resolve performance issues in the infrastructure whether they are caused by the network, application, or server. Learn more about EAA at <http://www.extremenetworks.com/product/extremeanalytics/>.

Extreme Management Center

Extreme Management Center (Management Center), formerly Netsight™, is a web-based control interface that provides centralized visibility into your network. Management Center reaches beyond ports, VLANs, and SSIDs and provides detailed control of individual users, applications, and protocols. When coupled with wireless and Identity & Access Management products, Management Center becomes the central location for monitoring and managing all the components in the infrastructure. Learn more about Management Center at <http://www.extremenetworks.com/product/management-center/>.

ExtremeCloud

ExtremeCloud is a cloud-based network management Software as a Service (SaaS) tool. ExtremeCloud allows you to manage users, wired and wireless devices, and applications on corporate and guest networks. You can control the user experience with smarter edges – including managing QoS, call admission control, secure access policies, rate limiting, multicast, filtering, and traffic forwarding, all from an intuitive web interface. Learn more about ExtremeCloud at <http://www.extremenetworks.com/product/extremecloud/>.

ExtremeControl

ExtremeControl, formerly Extreme Access Control™ (EAC), is a set of management software tools that use information gathered by a hardware engine to control policy to all devices on the network. The software allows you to automate and secure access for all devices on the network from a central dashboard, making it easier to roll out security and identity policies across the wired and wireless network. Learn more about ExtremeControl at <https://www.extremenetworks.com/product/extremecontrol/>.

ExtremeSwitching

ExtremeSwitching is the family of products comprising different switch types: **Modular** (X8 and 8000 series [formerly BlackDiamond] and S and K series switches); **Stackable** (X-series and A, B, C, and 7100 series switches); **Standalone** (SSA, X430, and D, 200, 800, and ISW series); and **Mobile Backhaul** (E4G). Learn more about ExtremeSwitching at <http://www.extremenetworks.com/products/switching-routing/>.

ExtremeWireless

ExtremeWireless products and solutions offer high-density WiFi access, connecting your organization with employees, partners, and customers everywhere they go. The family of wireless products and solutions includes APs, wireless appliances, and software. Learn more about ExtremeWireless at <http://www.extremenetworks.com/products/wireless/>.

ExtremeXOS

ExtremeXOS, a modular switch operating system, is designed from the ground up to meet the needs of large cloud and private data centers, service providers, converged enterprise edge networks, and everything in between. Based on a resilient architecture and protocols, ExtremeXOS supports network virtualization and standards-based SDN capabilities like VXLAN gateway, OpenFlow, and OpenStack Cloud orchestration. ExtremeXOS also supports comprehensive role-based policy. Learn more about ExtremeXOS at <http://www.extremenetworks.com/product/extremexos-network-operating-system/>.

FDB

The switch maintains a database of all MAC address received on all of its ports and uses this information to decide whether a frame should be forwarded or filtered. Each forwarding database (FDB) entry consists of the MAC address of the sending device, an identifier for the port on which the frame was received, and an identifier for the VLAN to which the device belongs. Frames destined for devices that are not currently in the FDB are flooded to all members of the VLAN. For some types of entries, you configure the time it takes for the specific entry to age out of the FDB.

FHSS

Frequency-Hopping Spread Spectrum is a transmission technology used in Local Area Wireless Network (LAWN) transmissions where the data signal is modulated with a narrowband carrier signal that 'hops' in a random but predictable sequence from frequency to frequency as a function of time over a wide band of frequencies. This technique reduces interference. If synchronized properly, a single logical channel is maintained. (Compare with [*DSSS \(Direct-Sequence Spread Spectrum\)*](#).)

GARP

Generic Attribute Registration Protocol registers an attribute with other participants. It is specified in IEEE 802.1D-2004, clause 12.

GVRP

GARP VLAN Registration Protocol is used to dynamically register VLANs on ports. It is specified in IEEE 802.1Q-2005, clause 11. GVRP is an example of the use of GARP, hence the G in GVRP.

IBSS

An IBSS is the 802.11 term for an ad hoc network. See [*ad hoc mode*](#).

ICMP

Internet Control Message Protocol is the part of the TCP/IP protocol that allows generation of error messages, test packets, and operating messages. For example, the ping command allows you to send ICMP echo messages to a remote IP device to test for connectivity. ICMP also supports traceroute, which identifies intermediate hops between a given source and destination.

IGMP

Hosts use Internet Group Management Protocol to inform local routers of their membership in multicast groups. Multicasting allows one computer on the Internet to send content to multiple other computers that have identified themselves as interested in receiving the originating computer's content. When all hosts leave a group, the router no longer forwards packets that arrive for the multicast group.

LACP

Link Aggregation Control Protocol is part of the IEEE 802.3ad and automatically configures multiple aggregated links between switches.

LAG

A Link Aggregation Group is the logical high-bandwidth link that results from grouping multiple network links in link aggregation (or load sharing). You can configure static LAGs or dynamic LAGs (using the LACP).

LLDP

Link Layer Discovery Protocol conforms to IEEE 802.1ab and is a neighbor discovery protocol. Each LLDP-enabled device transmits information to its neighbors, including chassis and port identification, system name and description, VLAN names, and other selected networking information. The protocol also specifies timing intervals in order to ensure current information is being transmitted and received.

MAC

Media Access Control layer. One of two sublayers that make up the Data Link Layer of the OSI model. The MAC layer is responsible for moving data packets to and from one NIC to another across a shared channel.

MD5

Message-Digest algorithm is a hash function that is commonly used to generate a 128-bit hash value. It was designed by Ron Rivest in 1991. MD5 is officially defined in RFC 1321 - The MD5 Message-Digest Algorithm.

MIB

Management Information Bases make up a database of information (for example, traffic statistics and port settings) that the switch makes available to network management systems. MIB names identify objects that can be managed in a network and contain information about the objects. MIBs provide a means to configure a network device and obtain network statistics gathered by the device. Standard, minimal MIBs have been defined, and vendors often have private enterprise MIBs.

MIC

Message Integrity Check (or Code), also called 'Michael', is part of WPA and TKIP. The MIC is an additional 8-byte code inserted before the standard 4-byte ICV appended in by standard WEP to the 802.11 message. This greatly increases the difficulty in carrying out forgery attacks. Both integrity check mechanisms are calculated by the receiver and compared against the values sent by the sender in the frame. If the values match, there is assurance that the message has not been tampered with.

MSTI

Multiple Spanning Tree Instances control the topology inside an MSTP region. An MSTI is a spanning tree domain that operates within a region and is bounded by that region; and MSTI does not exchange BPDUs or send notifications to other regions. You can map multiple VLANs to an MSTI; however, each VLAN can belong to only one MSTI. You can configure up to 64 MSTIs in an MSTP region.

MSTP

Multiple Spanning Tree Protocol, based on IEEE 802.1Q-2003 (formerly known as IEEE 892.1s), allows you to bundle multiple VLANs into one STP topology, which also provides enhanced loop protection and better scaling. MSTP uses RSTP as the converging algorithm and is compatible with legacy STP protocols.

netmask

A netmask is a string of 0s and 1s that mask, or screen out, the network part of an IP address, so that only the host computer part of the address remains. A frequently-used netmask is 255.255.255.0, used for a Class C subnet (one with up to 255 host computers). The ".0" in the netmask allows the specific host computer address to be visible.

PEAP

Protected Extensible Authentication Protocol is an IETF draft standard to authenticate wireless LAN clients without requiring them to have certificates. In PEAP authentication, first the user authenticates the authentication server, then the authentication server authenticates the user. If the first phase is successful, the user is then authenticated over the SSL tunnel created in phase one using EAP-Generic Token Card (EAP-GTC) or Microsoft Challenged Handshake Protocol Version 2 (MSCHAP V2). (See also [EAP-TLS/EAP-TTLS](#).)

PHY

The Physical Interface Transceiver is the device that implements the Ethernet physical layer (IEEE-802.3).

PoE

The Power over Ethernet standard (IEEE 802.3af) defines how power can be provided to network devices over existing Ethernet connections, eliminating the need for additional external power supplies.

QoS

Quality of Service is a technique that is used to manage network resources and guarantee a bandwidth relationship between individual applications or protocols. A communications network transports a multitude of applications and data, including high-quality video and delay-sensitive data such as real-time voice. Networks must provide secure, predictable, measurable, and sometimes guaranteed services. Achieving the required QoS becomes the secret to a successful end-to-end business solution.

RADIUS

RADIUS is a client/server protocol and software that enables remote access servers to communicate with a central server to authenticate dial-in users and authorize their access to the requested system or service. RADIUS allows a company to maintain user profiles in a central database that all remote servers can share. It provides better security, allowing a company to set up a policy that can be applied at a single administered network point. With RADIUS, you can track usage for billing and for keeping network statistics.

RFC

The IETF Request for Comments describe the definitions and parameters for networking. The RFCs are catalogued and maintained on the IETF RFC website: www.ietf.org/rfc.html.

RIP

This IGP vector-distance routing protocol is part of the TCP/IP suite and maintains tables of all known destinations and the number of hops required to reach each. Using Routing Information Protocol, routers periodically exchange entire routing tables. RIP is suitable for use only as an IGP.

root bridge

In STP, the root bridge is the bridge with the best bridge identifier selected to be the root bridge. The network has only one root bridge. The root bridge is the only bridge in the network that does not have a root port.

RSA

RSA is one of the first practicable public-key cryptosystems. It is widely used for secure data transmission. In such a cryptosystem, the encryption key is public and differs from the decryption key which is kept secret. In RSA, this asymmetry is based on the practical difficulty of factoring the product of two large prime numbers. This is called the factoring problem. RSA stands for Ron Rivest, Adi Shamir and Leonard Adleman, who first publicly described the algorithm in 1977. Clifford Cocks, an English mathematician, developed an equivalent system in 1973, but it wasn't declassified until 1997.

SMTP

Simple Mail Transfer Protocol uses the TCP to provide a mail service modeled on the FTP file transfer service. SMTP transfers mail between systems.

SNMP

Simple Network Management Protocol is a standard that uses a common software agent to remotely monitor and set network configuration and runtime parameters. SNMP operates in a multivendor environment, and the agent uses MIBs, which define what information is available from any manageable network device. You can also set traps using SNMP, which send notifications of network events to the system log.

SNTP

Simple Network Time Protocol is used to synchronize the system clocks throughout the network. An extension of NTP, SNTP can usually operate with a single server and allows for IPv6 addressing.

SSH

Secure Shell, sometimes known as Secure Socket Shell, is a UNIX-based command interface and protocol of securely gaining access to a remote computer. With SSH commands, both ends of the client/server connection are authenticated using a digital certificate, and passwords are protected by being encrypted. At Extreme Networks, the SSH is a separate software module, which must be downloaded separately. (SSH is bundled with SSL in the software module.)

SSL

Secure Socket Layer is a protocol for transmitting private documents using the Internet. SSL works by using a public key to encrypt data that is transferred over the SSL connection. SSL uses the public-and-private key encryption system, which includes the use of a digital certificate. SSL is used for other applications than SSH, for example, OpenFlow.

STP

Spanning Tree Protocol, defined in IEEE 802.1d, used to eliminate redundant data paths and to increase network efficiency. STP allows a network to have a topology that contains physical loops; it operates in bridges and switches. STP opens certain paths to create a tree topology, thereby preventing packets from looping endlessly on the network. To establish path redundancy, STP creates a tree that spans all of the switches in an extended network, forcing redundant paths into a standby, or blocked, state.

STP allows only one active path at a time between any two network devices (this prevents the loops) but establishes the redundant links as a backup if the initial link should fail. If STP costs change, or if one network segment in the STP becomes unreachable, the spanning tree algorithm reconfigures the STP topology and re-establishes the link by activating the standby path.

superloop

In EAPS, a superloop occurs if the common link between two EAPS domains goes down and the master nodes of both domains enter the failed state putting their respective secondary ports into the forwarding state. If there is a data VLAN spanning both EAPS domains, this action forms a loop between the EAPS domains.

syslog

A protocol used for the transmission of event notification messages across networks, originally developed on the University of California Berkeley Software Distribution (BSD) TCP/IP system implementations, and now embedded in many other operating systems and networked devices. A device generates a messages, a relay receives and forwards the messages, and a collector (a syslog server) receives the messages without relaying them.

syslog uses the UDP as its underlying transport layer mechanism. The UDP port that has been assigned to syslog is 514. (RFC 3164)

TACACS+

Terminal Access Controller Access Control System, often run on UNIX systems, provides access control for routers, network access servers, and other networked computing devices via one or more centralized servers. TACACS+ provides separate authentication, authorization, and accounting services. User passwords are administered in a central database rather than in individual routers, providing easily scalable network security solutions.

TCP

Transmission Control Protocol uses the Internet Protocol (IP) to exchange messages between computers. It is known as a connection-oriented protocol, which means that a connection is established and maintained until messages have been exchanged. TCP divides messages into packets for transfer via IP and reassembles the packets into complete messages for delivery to the receiving computer. Common network applications that use TCP include the worldwide web, email, and FTP.

Telnet

Teletype Network is a terminal emulation protocol that enables the Telnet client to control the Telnet server and communicate with other servers on the network. To start a Telnet session, users must enter a valid username and password. After logging in, they can enter commands as if they were working directly on the server console.

TLS

Transport Layer Security. See [SSL \(Secure Socket Layer\)](#).

TLV

An LLDP frame that can contain multiple pieces of information. Each piece is known as a Type Length Value.

UDP

User Datagram Protocol is an efficient but unreliable, connectionless protocol that is layered over IP (as is TCP). Application programs must supplement the protocol to provide error processing and retransmitting data. UDP is an OSI Layer 4 protocol.

USM

In SNMP version 3, the user-based security model uses the traditional concept of user names to associate with security levels to support secure network management.

VLAN

The term VLAN is used to refer to a collection of devices that communicate as if they are on the same physical LAN. Any set of ports (including all ports on the switch) is considered a VLAN. LAN segments are not restricted by the hardware that physically connects them. The segments are defined by flexible user groups you create with the CLI.

WLAN

Wireless Local Area Network.

Index

A

ACL 298

B

broadcast 97

C

CBC 140
Cipher Block-Chaining 140
conventions
 notice icons 6
 text 6

D

documentation
 feedback 7
 location 7, 8
dot1x port-control 341
dot1x system-auth-control 341

H

Hash 97

I

ip address 338
ip routing 338

L

listener, MSRP 251

M

Management Information Base 140
MD5 97
Message digest 5 97
MIB 140
MMRP:definition 251
MMRP:statistics 254
MRP:global settings 252
MRP:port settings 253
MSRP:definition 251
Multiple Registration Protocol 251

O

Object ID 140
OID 140
Open Source Declaration 7, 8

R

RADIUS accounting mode 341

S

Simple Network Time Protocol 97
SNTP 97
status HTML pages 19
support, *see* technical support

T

talker, MSRP 251
technical support
 contacting 7
Time levels 97

U

unicast 97
User Security Model 140
USM 140

V

VLAN database 338
VLAN participation 338
VLAN port tagging 338
VLAN routing 338

W

web interface: panel 17