LINKSYS™

User Guide

Managed Switch LGS5XX

Table of Contents

CHAPTER 1 Getting Started

Getting Started

This section provides an introduction to the Web-based configuration utility, and covers the following topics

- Configuring with the Console Port
- Starting the Web-based Configuration Utility
- Interface Naming Conventions
- Window Navigation

To configure with the Console Port, do the following:

1. Use a provided serial cable to connect to console port
2. Start a terminal application such as Hyper Terminal on your computer
3. Configure the utility with 11520 bit per second, 8 data bits, no parity, 1 stop bit and no flow control. (The firmware supports autobaud detection, the device will detect the speed after pressing Enter.)
4. Type in default user name: admin, and password: admin
5. Enter to access menu CLI

The following menu is displayed:

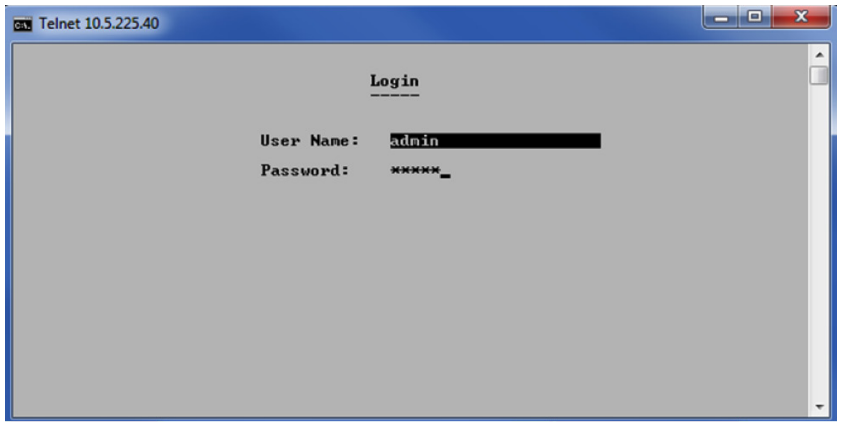

1. Enter your user name and password.

   The main menu is displayed:

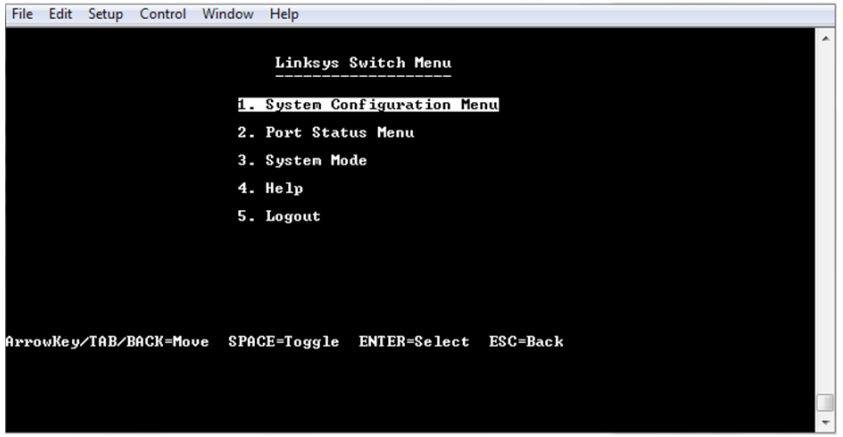

2. Continue configuring the device.
3. Click Logout to log out of the CLI menu.

Starting the Web-based Configuration Utility

This section describes how to navigate the Web-based switch configuration utility. If you are using a pop-up blocker, make sure it is disabled.

The following browsers are supported:

- Firefox (versions 16 and latest)
- IE version (versions 9, 10)

Browser Restrictions

If you are using IPv6 interfaces on your management station, use the IPv6 global address and not the IPv6 link local address to access the device from your browser.

Launching the Configuration Utility

To open the Web-based configuration utility, do the following:

STEP 1 Open a Web browser.

STEP 2 Enter the IP address of the device you are configuring in the address bar on the browser, and then press Enter.

NOTE:
When the device is using the factory default IP address of 192.168.1.251, its power LED flashes continuously. When the device is using a DHCP assigned IP address or an administrator-configured static IP address, the power LED is on solid.

Logging In

The default username is **admin** and the default password is **admin**.

To log in to the Web-based GUI:

STEP 1 Open the GUI. The Login page is displayed.

STEP 2 Enter the username/password. The password can contain up to 64 ASCII characters.

Logging Out

By default, the application logs out after ten minutes of inactivity.

CAUTION:
Unless the Running Configuration is copied to the Startup Configuration, rebooting the device will remove all changes made since the last time the file was saved. Save the Running Configuration to the Startup Configuration before logging off to preserve any changes you made during this session.

When you click "*Quick Start > Save Your Configurations*, the Configuration File Copy page appears. Save the Running

Configuration file by copying it to the Startup Configuration file.

To log out, click Logout in the top right corner of any page. The system logs out of the device.

When a timeout occurs or you intentionally log out of the system, a message appears and the Login page appears, with a message indicating the logged-out state.

Interface Naming Conventions

Within the GUI, interfaces are denoted by linking the following elements:

- Interface Number: Port, LAG or VLAN ID

Window Navigation

This section describes the features of the Web-based switch configuration utility.

Application Header

The Application Header appears on every page. It provides the following application links:

Application Link Name	Description
Logout	Click to log out of the Web-based switch configuration utility.
Firmware Version	Display the device version number.
Help	Click for the link to this administration guide.

Management Buttons

The following table describes the commonly used buttons that appear on various pages in the system.

Button Name	Description
Add	Click to display the related Add page and add an entry to a table. Enter the information and click Apply to save it to the Running Configuration. Click Close to return to the main page. Click Save to display the Configuration File Copy page and save the Running Configuration to the Startup Configuration file type on the device.
Apply	Click to apply changes to the Running Configuration on the device. If the device is rebooted, the Running Configuration is lost unless it is saved to the Startup Configuration file type or another file type. Click Save to display the Configuration File Copy page and save the Running Configuration to the Startup Configuration file type on the device.
Close	Click to return to the previous page. Any changes not applied are cleared.
Clear All	Click to clear the statistic counters for all interfaces.
Clear	Click to clear information, such a counters of an interface or all interface, or log files.
Delete	After selecting an entry in the table, click Delete to remove.
Edit	Select the entry and click Edit. The Edit page appears, and the entry can be modified. 1. Click Apply to save the changes to the Running Configuration. 2. Click Close to return to the main page.
Search	Enter the query filtering criteria and click Search. The results are displayed on the page.
Refresh	Click Refresh to refresh the counter values.
Test or Start	Click Test to perform the related tests.
View or View All	Click View to display details associated with the entry selected or for all entries (respectively).

CHAPTER 2 System Status

System Status

This section describes how to view device statistics. It covers the following topics:

- System Summary
- RMON
- Interface Statistics

System Summary

The System Summary page provides a graphic view of the device, and displays device status, hardware information, firmware version information, general PoE status, and other items.

To view system information, click System Status > System Summary. The System Summary page contains system and hardware information.

- System Mode—Specifies whether the system is operating in Layer 2 or Layer 3 system mode.
- System Description—A description of the system.
- System Location—Physical location of the device. Click Edit to go the System Information page to enter this value.
- System Contact—Name of a contact person. Click Edit to go the System Information page to enter this value.
- Host Name—Name of the device. By default, the device host name is composed of the name of the switch followed by the final six digits in the device's MAC address.
- Base MAC Address—Device MAC address.
- SNMP Object ID— The unique vendor identification of the network management subsystem assigned by Internet Assigned Numbers Authority.
- Firmware Version—Firmware version number.
- Boot Code Version—Boot version number.
- Hardware Version —Hardware version number of the device.
- Serial Number—Serial number.

Device Status

- Fan Status—Applicable only to models that have fans. The following values are possible:
- OK—Fan is operating normally.
- Fail—Fan is not operating correctly.
- Date & Time—System date and time.
- System Uptime—Length of time since last reboot.

RMON

The Statistics page displays detailed information regarding packet sizes and information regarding physical layer errors. The information displayed is according to the RMON (Remote Network Monitoring) standard. An oversized packet is defined as an Ethernet frame with the following criteria:

- Packet length is greater than MRU byte size.
- Collision event has not been detected.
- Late collision event has not been detected.
- Received (Rx) error event has not been detected.
- Packet has a valid CRC.

To view RMON statistics and/or set the refresh rate, do the following:

STEP 1 Click System Status > RMON > Statistics.

STEP 2 Select the Interface for which statistics are to be displayed.

STEP 3 Select the Refresh Rate, the time period that passes before the interface statistics are refreshed.

The statistics are displayed for the selected interface.

- Bytes Received—Number of octets received, including bad packets and FCS octets, but excluding framing bits.
- Drop Events—Number of packets dropped.
- Packets Received—Number of good packets received, including Multicast and Broadcast packets.
- Broadcast Packets Received—Number of good Broadcast packets received. This number does not include Multicast packets.
- Multicast Packets Received—Number of good Multicast packets received.
- CRC & Align Errors—Number of CRC and Align errors that have occurred.
- Undersize Packets—Number of undersized packets (less than 64 octets) received.
- Oversize Packets—Number of oversized packets (over 2000 octets) received.
- Fragments—Number of fragments (packets with less than 64 octets, excluding framing bits, but including Frame Check Sequence octets) received.
- Jabbers—Total number received packets that were longer than 1632 octets. This number excludes frame bits, but includes FCS octets that had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a non-integral octet (Alignment Error) number. A jabber packet is defined as an Ethernet frame that satisfies the following criteria:
- Packet data length is greater than MRU.
- Packet has an invalid CRC.
- Received (Rx) Error Event has not been detected.
- Collisions—Number of collisions received. If Jumbo Frames are enabled, the threshold of Jabber Frames is raised to the maximum size of Jumbo Frames.
- Frames of 64 Bytes—Number of frames, containing 64 bytes that were received.
- Frames of 65 to 127 Bytes—Number of frames, containing 65-127 bytes that were received.
- Frames of 128 to 255 Bytes—Number of frames, containing 128-255 bytes that were received.
- Frames of 256 to 511 Bytes—Number of frames, containing 256-511 bytes that were received.
- Frames of 512 to 1023 Bytes—Number of frames, containing 512-1023 bytes that were received.
- Packets of 1024 and More Bytes—Number of frames, containing 1024-2000 bytes, and Jumbo Frames, that were received.

To clear or view statistics counters, do the following:

- Click Refresh to refresh the counters on the page.
- Click Clear to clear the selected interfaces counters.
- Click View All to see all ports on a single page.

Interface Statistics

The Interface Statistics page displays traffic statistics per port. The refresh rate of the information can be selected.

This page is useful for analyzing the amount of traffic that is both sent and received and its dispersion (Unicast, Multicast, and Broadcast).

To display Ethernet statistics and/or set the refresh rate, do the following:

STEP 1 Click System Status > Interface Statistics.

STEP 2 Enter the parameters.

- Interface—Select the specific interface for which Ethernet statistics are to be displayed.
- Refresh Rate—Select the time period that passes before the interface Ethernet statistics are refreshed. The available options are as follows:
- No Refresh—Statistics are not refreshed.
- 15 Sec—Statistics are refreshed every 15 seconds.
- 30 Sec—Statistics are refreshed every 30 seconds.
- 60 Sec—Statistics are refreshed every 60 seconds.

The Receive Statistics area displays information about incoming packets.

- Total Octets—Octets received, including bad packets and FCS octets, but excluding framing bits.
- Unicast Packets—Good Unicast packets received.
- Multicast Packets—Good Multicast packets received.
- Broadcast Packets—Good Broadcast packets received.
- Error Packets—Packets with errors received.

The Transmit Statistics area displays information about outgoing packets.

- Total Octets—Octets transmitted, including bad packets and FCS octets, but excluding framing bits.
- Unicast Packets—Good Unicast packets transmitted.
- Multicast Packets—Good Multicast packets transmitted.
- Broadcast Packets—Good Broadcast packets transmitted.

To clear or view statistics counters, do the following:

- Click Refresh to refresh the counters on the page.
- Click Clear to clear the selected interfaces counters.
- Click View All to see all ports on a single page.

Chapter 3 Quick Start

Quick Start

This section describes how to view device statistics.

To simplify device configuration through quick navigation, the Quick Start page provides links to the most commonly used pages.

Link Name (on the Page)	Linked Page
Configure User Accounts and Management Access	User Access & Accounts
Configure Device IP Address	IPv4 Interface
Create VLANs	VLANs
Configure VLAN Memberships	VLAN Memberships
Save Your Configuration	Configuration File Copy

Clicking on the Support link takes you to the device product support page.

Chapter 4 System Management

System Management

This chapter describes the following topics:

- System Information
- Time
- SNMP
- Logs

System Information

To enter system information, do the following:

STEP 1 Click Configuration > System Management > System Information.

STEP 2 View or modify the system settings.

- System Description—Displays a description of the device.
- System Location—Enter the location where the device is physically located.
- System Contact—Enter the name of a contact person.
- System Host Name—Select the host name of this device, which used in the prompt of CLI commands.
- Default—The default host name (System Name) of these switches is switch123456, where 123456 represents the last three bytes of the device MAC address in hex format.
- User Defined—Enter the host name. Use only letters, digits, and hyphens. Host names cannot begin or end with a hyphen. No other symbols, punctuation characters, or blank spaces are permitted (as specified in RFC1033, 1034, 1035).

STEP 3 Click Apply to save the values in the Running Configuration file.

Time

This section describes the options for configuring the system time, time zone, and Daylight Savings Time (DST). It covers the following topics:

- Overview
- System Time
- SNTP Unicast Server

Overview

Synchronized system clocks provide a frame of reference between all devices on the network. Network time synchronization is critical because every aspect of managing, securing, planning, and debugging a network involves determining when events occur. Without synchronized clocks, accurately correlating log files between devices when tracking security breaches or network usage is impossible.

Synchronized time also reduces confusion in shared file systems, as it is important for the modification times to be consistent, regardless of the machine on which the file systems reside.

For these reasons, it is important that the time configured on all of the devices on the network is accurate.

NOTE:
The device supports Simple Network Time Protocol (SNTP) and when enabled, the device dynamically synchronizes the device time with time from an SNTP server. The device operates only as an SNTP client, and cannot provide time services to other devices.

System time can be set manually by the user or dynamically from an SNTP server. If an SNTP server is chosen, the manual time settings are overwritten when communications with the server are established.

As part of the boot process, the device always configures the time, time zone, and DST. These parameters are obtained from SNTP, values set manually, or, if all else fails, from the factory defaults.

Time Configuration Methods

The following methods are available for setting the system time on the device:

- Manual—You must manually set the time.
- SNTP—Time can be received from SNTP time servers. SNTP ensures accurate network time synchronization of the device up to the millisecond by using an SNTP server for the clock source. When specifying an SNTP server, if choosing to identify it by host name, three suggestions are given in the GUI:
- time-a.timefreq.bldrdoc.gov
- time-b.timefreq.bldrdoc.gov
- time-c.timefreq.bldrdoc.gov

After the time has been set by any of the above sources, it is not set again by the browser.

NOTE:
SNTP is the recommended method for time setting.

Time Zone and Daylight Savings Time (DST)

The Time Zone and DST can be set on the device in the following ways:

- Dynamic configuration of the device through a DHCP server, as follows:
- Dynamic DST, when enabled and available, always takes precedence over the manual configuration of DST.
- If the server supplying the source parameters fails, or dynamic configuration is disabled by the user, the manual settings are used.
- Dynamic configuration of the time zone and DST continues after the IP address lease time has expired.
- Manual configuration of the time zone and DST becomes the Operational time zone and DST, only if the dynamic configuration is disabled or fails.

NOTE:
The DHCP server must supply DHCP option 100 in order for dynamic time zone configuration to take place.

SNTP Modes

The device can receive the system time from an SNTP server in one of the following ways:

- Client Multicast/Anycast Reception (passive mode) SNTP servers broadcast the time, and the device listens to these Multicast/Anycast transmissions. When the device is in this mode, there is no need to define a Unicast SNTP server.
- Client Multicast/Anycast Transmission (active mode)—The device, as an SNTP client, periodically requests SNTP time updates. This mode works in either of the following ways:
- SNTP Anycast Client Mode—The device broadcasts time request packets to all SNTP servers in the subnet, and waits for a response.
- Unicast SNTP Server Mode—The device sends Unicast queries to a list of manually-configured SNTP servers, and waits for a response. The device supports having all of the above modes active at the same time and selects the best system time received from an SNTP server, according to an algorithm based on the closest stratum (distance from the reference clock).

System Time

Use the System Time page to select the system time source. If the source is manual, you can enter the time here.

CAUTION:
If the system time is set manually and the device is rebooted, the manual time settings must be reentered.

To define system time, do the following:

STEP 1 Click Configuration > System Management > Time > System Time.

The current time on the device is displayed. This shows the DHCP time zone or the acronym for the user-defined time zone if these were defined.

STEP 2 Enter these parameters:

Clock Source—Select the source used to set the system clock.

- SNTP—If you enable this, the system time is obtained from an SNTP server. To use this feature, you must also configure a connection to an SNTP server in the SNTP Unicast Server page.

- SNTP Client Unicast—Select to enable client Unicast mode.
- SNTP IPv4 Multicast Rx—Select to receive SNTP IPv4 Multicast synchronization packets requesting system time information. The packets are transmitted to all SNTP servers on the subnet.
- SNTP IPv4 Anycast Tx—Select to transmit SNTP IPv4 Anycast synchronization packets requesting system time information. The packets are transmitted to all SNTP servers on the subnet.
- SNTP IPv6 Multicast Rx—Select to receive SNTP IPv6 Multicast synchronization packets requesting system time information. The packets are transmitted to all SNTP servers on the subnet.
- SNTP IPv6 Anycast Tx—Select to transmit SNTP IPv6 Anycast synchronization packets requesting system time information. The packets are transmitted to all SNTP servers on the subnet.
- Manual Date/Time—Set the date and time manually. The local time is used when there is no alternate source of time, such as an SNTP server.

Time Zone—The local time is used via the DHCP server or Time Zone offset.

- Time Zone from DHCP—Select to enable dynamic configuration of the time zone and the DST from the DHCP server. Whether one or both of these parameters can be configured depends on the information found in the DHCP packet. If this option is enabled, you must also enable DHCP client on the device. The DHCP Client supports Option 100 providing dynamic time zone setting.
- DHCP Time Zone—Displays the acronym of the time zone configured from the DHCP server. This acronym appears in the Actual Time field.
- Time Zone Offset—Select the difference in hours between Greenwich Mean Time (GMT) and the local time. For example, the Time Zone Offset for Paris is GMT +1, while the Time Zone Offset for New York is GMT – 5.
- Time Zone Acronym—Enter a user-defined name that represents the time zone you have configured. This acronym appears in the Actual Time field.
- Daylight Savings Time—Select how DST is defined:
- Daylight Savings — Select to enable Daylight Saving Time.
- Time Set Offset—Enter the number of minutes offset from GMT ranging from1—1440. The default is 60.
- Daylight Savings Type—Click one of the following:
- USA — DST is set according to the dates used in the USA.
- European — DST is set according to the dates used by the European
- Union and other countries that use this standard.
- By Dates — DST is set manually, typically for a country other than the USA or a European country. Enter the following parameters:
- Selecting By Dates allows customization of the start and stop of DST:
- From — Day and time that DST starts.
- To — Day and time that DST ends.
- Recurring — DST occurs on the same date every year.

Selecting Recurring allows different customization of the start and stop of DST:

- From — Date when DST begins each year.
- Day — Day of the week on which DST begins every year.
- Week — Week within the month from which DST begins every year.
- Month — Month of the year in which DST begins every year.
- Time — The time at which DST begins every year.
- To — Date when DST ends each year. For example, DST ends locally every fourth Friday in October at 5:00 am. The parameters are as follows:
- Day — Day of the week on which DST ends every year.
- Week — Week within the month from which DST ends every year.
- Month—Month of the year in which DST ends every year.
- Time—The time at which DST ends every year.

STEP 3 Click Apply. The system time values are written to the Running Configuration file.

SNTP Unicast Server

Up to 16 Unicast SNTP servers can be configured.

NOTE:
To specify a Unicast SNTP server by name, you must first configure DNS server(s) on the device (see DNS Settings). To add a Unicast SNTP server, SNTP Client Unicast must be enabled.

To add a Unicast SNTP server:

STEP 1 Click Configuration > Time > SNTP Unicast Server.

This page displays the following information for each configured Unicast SNTP server:

- SNTP Server — SNTP server IP address. The preferred server, or host name, is chosen according to its stratum level.
- SNTP Server Status—SNTP server status. The possible values are as follows:
- Up — SNTP server is currently operating normally.
- Down — SNTP server is currently not available.
- Unknown — SNTP server is currently being searched for by the device.
- In Process — Occurs when the SNTP server has not fully trusted its own time server (i.e. when first booting up the SNTP server).
- Stratum Level—Distance from the reference clock expressed as a numerical value. An SNTP server cannot be the primary server (stratum level 1) unless polling interval is enabled.
- Offset—The estimated offset of the server's clock relative to the local clock, in milliseconds. The host determines the value of this offset using the algorithm described in RFC 2030.
- Delay—The estimated round-trip delay of the server's clock relative to the local clock over the network path between them, in milliseconds. The host determines the value of this delay using the algorithm described in RFC 2030.
- Poll Interval—Displays whether polling is enabled or disabled.
- Last Response Time—Date and time of the last time a response was received from this SNTP server.

STEP 2 Click Add.

STEP 3 Enter the following parameters:

- SNTP Server—Select if the SNTP server is going to be identified by its IP address or if you are going to select a well-known SNTP server by name from the list.

NOTE:
To specify a well-known SNTP server, the device must be connected to the Internet and configured with a DNS server or configured so that a DNS server is identified by using DHCP. (See DNS Settings)

- IP Version—Select the version of the IP address: Version 4 or Version 6.
- IPv6 Address Type—Select the IPv6 address type (if IPv6 is used). The options are
- Global—The IPv6 address is a global Unicast IPV6 type that is visible and reachable from other networks.
- Link Local—The IPv6 address uniquely identifies hosts on a single network link. A link local address has a prefix of FE80, is not routable, and can be used for communication only on the local network. Only one link local address is supported. If a link local address exists on the interface, this entry replaces the address in the configuration.
- Interface—Select the link local interface (if IPv6 Address Type Link Local is selected) from the list.
- SNTP Server IP Address—Enter the SNTP server IP address. The format depends on which address type was selected.
- SNTP Server Name—Select the name of the SNTP server from a list of well-known NTP servers. If other is chosen, enter name of SNTP server in the adjacent field.
- Poll Interval—Select to enable polling of the SNTP server for system time information. All NTP servers that are registered for polling are polled, and the clock is selected from the server with the lowest stratum level (distance from the reference clock) that is reachable. The server with the lowest stratum is considered to be the primary server. The server with the next lowest stratum is a secondary server, and so forth. If the primary server is down, the device polls all servers with the polling setting enabled, and selects a new primary server with the lowest stratum.

STEP 4 Click Apply. The STNP server is added, and you are returned to the main page.

SNTP Multicast/Anycast Interface

The device can be in active and/or passive mode (see SNTP Modes for more information).

To enable receiving SNTP packets from all servers on the subnet and/or to enable transmitting time requests to SNTP servers, do the following:

STEP 1 Click Administration > Time Settings > SNTP Multicast/Anycast.

STEP 2 If the system is in Layer 3 system mode, click Add to select the interface for SNTP reception/transmission. Select an interface and select the reception/transmission options.

STEP 3 Click Apply to save the settings to the Running Configuration file.

SNMP

This section describes the Simple Network Management Protocol (SNMP) feature that provides a method for managing network devices. SNMP v1/v2 are supported. It covers the following topics:

- Overview
- Feature Configuration

Overview

The following are the device model Object IDs (OIDs).

Mode Name	Description	Object ID
LGS528	28-Port Managed Gigabit Switch	enterprises(1).linksys(3955).smb(1000).5.28.1
LGS552	52-Port Managed Gigabit Switch with two 10 Gigabit uplink	enterprises(1).linksys(3955).smb(1000).5.52.1
LGS528P	28-Port Managed PoE+ Gigabit Switch	enterprises(1).linksys(3955).smb(1000).5.28.1
LGS552P	52-Port Managed PoE+ Gigabit Switch with two10 Gigabit uplink	enterprises(1).linksys(3955).smb(1000).5.52.2

The private Object IDs are placed under enterprises(1).linksys(3955).smb(1000).switch01(201).

Feature Configuration

To enable SNMP do the following:

STEP 1 Click Configuration > System Management > Feature Configuration.

The following fields are displayed for each community:

- SNMP Management Station—Management station IP address that can access the SNMP community.
- Community—Password for this community.
- Access Mode—Access rights of the community. The options are as follows:
- Read Only—Management access is restricted to read-only. Changes cannot be made to the community.
- Read-Write—Management access is read-write. Changes can be made to the device configuration, but not to the community.

Logs

This section describes the Logs feature, which enables the device to generate multiple independent logs. It covers the following topics:

- Overview
- Log Management
- Remote Log Servers
- RAM Log
- Flash Memory Log

Overview

Each log is a set of messages describing system events. The device generates the following local logs:

- Log sent to the console interface.
- Log written into a cyclical list of logged events in the RAM and erased when the device reboots.
- Log written to a cyclical log-file saved to the Flash memory and persists across reboots.

In addition, you can send messages to remote SYSLOG servers in the form of SNMP traps and SYSLOG messages.

You can configure the messages that are written to each log by severity, and a message can go to more.

Log Management

You can enable or disable logging on the Log Management page.

You can select the events by severity level. Each log message has a severity level marked with the first letter of the severity level separated by dashes (-) on each side (except for Emergency that is indicated by the letter F). For example, the log message "%INIT-I-InitCompleted: … " has a severity level of I, meaning Informational.

The event severity levels are listed from the highest severity to the lowest severity, as follows:

- Emergency—System is not usable.
- Alert—Action is needed.
- Critical—System is in a critical condition.
- Error—System is in error condition.
- Warning—System warning has occurred.
- Notice—System is functioning properly, but a system notice has occurred.
- Informational—Device information.
- Debug—Detailed information about an event.

You can select different severity levels for RAM and Flash logs. These logs are displayed in the RAM Log page and Flash Memory Log page, respectively.

Selecting a severity level to be stored in a log causes all of the higher severity events to be automatically stored in the log. Lower severity events are not stored in the log.

For example, if Warning is selected, all severity levels that are Warning and higher are stored in the log (Emergency, Alert, Critical, Error, and Warning). No events with severity level below Warning are stored (Notice, Informational, and Debug).

To set global log parameters, do the following:

STEP 1 Click Configuration > System Management > Logs > Log Management.

STEP 2 Enter the parameters.

System Log

- Logging—Select to enable message logging.
- Originator Identifier—Enables adding an origin identifier to SYSLOG messages. The options are as follows:
- None—Do not include the origin identifier in SYSLOG messages.
- Hostname—Include the system hostname in SYSLOG messages.
- IPv4 Address—Include the IPv4 address of the sending interface in SYSLOG messages.
- IPv6 Address—Include the IPv6 address of the sending interface in
- SYSLOG messages.
- User Defined—Enter a description to be included in SYSLOG messages.

Log Settings

- Severity—Select the severity levels of the messages to be logged to the following:
- RAM Memory Logging—Severity levels of the messages to be logged to the RAM.
- Flash Memory Logging—Severity levels of the messages to be logged to the Flash memory.

STEP 3 Click Apply. The Running Configuration file is updated.

Remote Log Servers

The Remote Log Servers page enables defining remote SYSLOG servers where log messages are sent (using the SYSLOG protocol). For each server, you can configure the severity of the messages that it receives.

To define SYSLOG servers, do the following:

STEP 1 Click Configuration > System Management > Logs > Remote Log Servers.

The list of configured remote log servers is displayed.

STEP 2 Click Add.

STEP 3 Enter the parameters.

- Remote Log Server—Select whether to identify the remote log server by IP address or name.
- IP Version—Select the supported IP version.
- IPv6 Address Type—Select the IPv6 address type (if IPv6 is used). The options are as follows:
- Global—The IPv6 address is a global Unicast IPV6 type that is visible and reachable from other networks.
- Link Local—The IPv6 address uniquely identifies hosts on a single network link. A link local address has a prefix of FE80, is not routable, and can be used for communication only on the local network. Only one link local address is supported. If a link local address exists on the interface, this entry replaces the address in the configuration.
- Interface—Select the link local interface (if IPv6 Address Type Link Local is selected) from the list.
- Log Server IP Address—Enter the IP address of the log server if it is to be identified by address.
- Log Server Name—Enter the domain name of the log server if it is to be identified by name.

Server Settings

- UDP Port—Enter the UDP port to which the log messages are sent.
- Facility—Select a facility value from which system logs are sent to the remote server. Only one facility value can be assigned to a server. If a second facility code is assigned, the first facility value is overridden.
- Description—Enter a server description.
- Minimum Logging Level—Select the minimum level of system log messages to be sent to the server.

STEP 4 Click Apply. The SYSLOG server is added, and the Running Configuration file is updated.

RAM Log

The RAM Log page displays all messages that were saved in the RAM (cache) in chronological order. Entries are stored in the RAM log according to the configuration in the Log Management page.

To view log entries, click Configuration > System Management > Logs > RAM Log.

This page contains the following fields:

- Log Index—Log entry number.
- Log Time—Time when message was generated.
- Severity—Event severity.
- Description—Message text describing the event.

To clear the log messages, click Clear. The messages are cleared.

Flash Memory Log

The Flash Memory Log page displays the messages that were stored in the Flash memory, in chronological order. The minimum severity for logging is configured in the Log Management page. Flash logs remain when the device is rebooted. You can clear the logs manually.

To view the Flash logs, click Configuration > System Management > Logs >

Flash Memory Log.

This page contains the following fields:

- Log Index—Log entry number.
- Log Time—Time when message was generated.
- Severity—Event severity.
- Description—Message text describing the event.

To clear the messages, click Clear. The messages are cleared.

TCAM Resources

The TCAM Resources page is only displayed in Layer 3 mode.

TCAM holds the rules produced by applications, such as Access Control Lists (ACLs), Quality of Service (QoS), IP Routing and user-created rules.

Some applications reserve TCAM resources that will be required upon their initiation. Additionally, processes that initialize during system boot might configure some rules during the startup process.

To configure and view TCAM utilization:

STEP 1 Click Configuration > System Management > TCAM Resources.

The list of configured remote log servers will be displayed.

STEP 2 Enter the following fields:

- Maximum IPv4 TCAM Entries—Maximum TCAM entries available for IPv4 routing. Select one of the following options:
  - Use Default—Use the system value for this field.
  - User Defined—Enter the maximum number of TCAM entries that you determine will be used for IPv4 routing.

The following counters are displayed for TCAM utilization:

- IPv4 Hosts:
  - Count—Number of IPv4 interfaces configured on the switch.
  - TCAM Entries—Number of TCAM entries currently used by the known IPv4 nodes.
- IPv4 Interfaces:
  - Number of IPv4 interfaces configured on the switch.
  - TCAM Entries—Number of TCAM entries used by the configured IPv4 interfaces.
- IPv4 Routes.
  - Count—Number of known IP routes on the switch.
  - TCAM Entries—Number of TCAM entries currently used by the known IP routes.
- Total—Total number of TCAM entries.

The following counters are displayed for Non-IP TCAM Usage:

- Non-IP
  - In Use—Number of TCAM entries currently used by applications and features, excluding IP routing.
  - Maximum Allocated—Number of available TCAM entries that can be used by applications and features, excluding IP routing.

Chapter 5 Port Management

Port Management

This section describes port configuration, link aggregation, and the Green Ethernet feature.

It covers the following topics:

- Ports
- Link Aggregation
- Green Ethernet
- PoE

Ports

Workflow

To configure ports, perform the following actions:

1. Configure port by using the Ports page.
2. Enable/disable the Link Aggregation Control protocol, and configure the potential member ports to the desired LAGs by using the LAG Aggregation page. By default, all LAGs are empty.
3. Configure the Ethernet parameters, such as speed and auto-negotiation for the LAGs by using the LAGs page.
4. Configure the LACP parameters for the ports that are members or candidates of a dynamic LAG by using the LAGs page.
5. Configure Green Ethernet and 802.3 Energy Efficient Ethernet by using the Green Ethernet page.
6. Configure Green Ethernet energy mode and 802.3 Energy Efficient Ethernet per port by using the Green Ethernet page.
7. If PoE is supported and enabled for the device, configure the device as described in PoE.

Configure Port Settings

To configure port settings:

STEP 1 Click Configuration > Port Management > Ports.

STEP 2 Select Enable to support jumbo packets of up to 10 KB in size. If Jumbo Frames is not enabled (default), the system supports packet size up to 2,000 bytes. For Jumbo Frames to take effect, the device must be rebooted after the feature is enabled.

STEP 3 To update the port settings, select the desired port, and click Edit.

STEP 4 Modify the following parameters:

- Port—Select the port number.

Port Settings

- Operational Status—Displays whether the port is currently up or down. If the port is down because of an error, the description of the error is displayed.
- Administrative Mode—Select to bring the port up or down.
- Suspended Port—Select to reactivate a port that has been suspended. The reactivate operation brings the port up without regard to why the port was suspended.
- Protected Port—Select to make this a protected port. (A protected port is also referred to as a Private VLAN Edge (PVE).) The features of a protected port are as follows:
- Protected Ports provide Layer 2 isolation between interfaces (Ethernet ports and LAGs) that share the same VLAN.
- Packets received from protected ports can be forwarded only to unprotected egress ports. Protected port filtering rules are also applied to packets that are forwarded by software, such as snooping applications.

- Port protection is not subject to VLAN membership. Devices connected to protected ports are not allowed to communicate with each other, even if they are members of the same VLAN.
- Both ports and LAGs can be defined as protected or unprotected.
- Protected LAGs are described in the LAGs section.
- Auto Negotiation—Select to enable auto-negotiation on the port. Auto-negotiation enables a port to advertise its transmission speed, duplex mode, and flow control abilities to the port link partner.
- Port Speed—Configure the speed of the port. The port type determines the available speeds. You can designate this field only when port auto-negotiation is disabled.
- Duplex Mode—Select the port duplex mode. This field is configurable only when auto-negotiation is disabled, and the port speed is set to 10M or 100M. At port speed of 1G, the mode is always full duplex. The possible options are:
- Half—The interface supports transmission between the device and the client in only one direction at a time.
- Full—The interface supports transmission between the device and the client in both directions simultaneously.
- Auto Advertisement—Select the capabilities advertised by auto-negotiation when it is enabled. The optionsare as follows:
- Max Capability—All port speeds and duplex mode settings can be accepted.
- 10 Full Duplex—10 Mbps speed and Full Duplex mode.
- 10 Half Duplex—10 Mbps speed and Half Duplex mode.
- 100 Full Duplex—100 Mbps speed and Full Duplex mode.
- 100 Half Duplex—100 Mbps speed and Half Duplex mode.
- 1000 Full Duplex—1000 Mbps speed and Full Duplex mode.
- Back Pressure—Select the Back Pressure mode on the port (used with Half Duplex mode) to slow down the packet reception speed when the device is congested. It disables the remote port, preventing it from sending packets by jamming the signal.
- Flow Control—Enable or disable 802.3x Flow Control, or enable the auto-negotiation of flow control on the port (only when in Full Duplex mode).
- MDI/MDIX—the Media Dependent Interface (MDI)/Media Dependent
- Interface with Crossover (MDIX) status on the port. The options are as follows:
- MDIX—Select to swap the port's transmit and receives pairs.
- MDI—Select to connect this device to a station by using a straight-through cable.
- Auto—Select to configure this device to automatically detect the correct pinouts for the connection to another device.
- Description—Enter the port description.

STEP 5 Click Apply. The port settings are written to the Running Configuration file.

Link Aggregation

- This section describes how to configure LAGs. It covers the following topics:
- Overview
- LAGs

Overview

Link Aggregation Control Protocol (LACP) is part of the IEEE specification (802.3ad) that enables you to bundle several physical ports together to form a single logical channel (LAG). LAGs multiply the bandwidth, increase port flexibility, and provide link redundancy between two devices.

- Two types of LAGs are supported:
- Static—A LAG is static if the LACP is disabled on it. The ports assigned to a static LAG are always active members. After a LAG is manually created, the LACP option cannot be added or removed, until the LAG is edited and a member is removed (which can be added prior to applying), then the LACP button becomes available for editing.
- Dynamic—A LAG is dynamic if LACP is enabled on it. The ports assigned to dynamic LAG are candidate ports. LACP determines which candidate ports are active member ports. The non-active candidate ports are standby ports ready to replace any failing active member ports.

Load Balancing

- Traffic forwarded to a LAG is load-balanced across the active member ports, thus achieving an effective bandwidth close to the aggregate bandwidth of all the active member ports of the LAG.
- Traffic load balancing over the active member ports of a LAG is managed by a hash-based distribution function that distributes Unicast and Multicast traffic based on Layer 2 or Layer 3 packet header information.
- The device supports two modes of load balancing:
- By MAC Addresses—(Default) Based on the destination and source MAC
- addresses of all packets.
- By IP and MAC Addresses—Based on the destination and source IP addresses for IP packets, and destination and source MAC addresses for non-IP packets.

LAG Management

In general, a LAG is treated by the system as a single logical port. In particular, the LAG has port attributes similar to a regular port, such as state and speed.

The device supports four LAGs.

Every LAG has the following characteristics:

- All ports in a LAG must be of the same media type.
- To add a port to the LAG, it cannot belong to any VLAN except the default VLAN.
- Ports in a LAG must not be assigned to another LAG.
- No more than eight ports are assigned to a static LAG and no more than 16 ports can be candidates for a dynamic LAG.
- All the ports in a LAG must have auto-negotiation disabled, although the LAG can have auto-negotiation enabled.
- When a port is added to a LAG, the configuration of the LAG is applied to the port. When the port is removed from the LAG, its original configuration is reapplied.
- Protocols, such as Spanning Tree, consider all the ports in the LAG to be one port.

Default Settings and Configuration

Ports are not members of a LAG and are not candidates to become part of a LAG.

Static and Dynamic LAG Workflow

After a LAG has been manually created, LACP cannot be added or removed until the LAG is edited and a member is removed. Only then the LACP field is activated.

- To configure a static LAG, perform the following actions:

1. Disable LACP on the LAG to make it static. Assign up to eight member ports to the static LAG in the Port List to the LAG Port Member list. Perform these actions in the LAGs page.
2. Configure various aspects of the LAG, such as speed and flow control by using the Edit LAG page.

To configure a dynamic LAG, perform the following actions:

1. Enable LACP on the LAG. Assign up to 16 candidates ports to the dynamic LAG by selecting and moving the ports from the Port List to the LAG Port Member List by using the LAGs page.
2. Configure various aspects of the LAG, such as speed and flow control by using the LAGs page.

LAGs

The LAGs page displays global and per-LAG settings. The page also enables you to configure the global settings and to select and edit the desired LAG on the Edit LAG Membership page.

- To define the member or candidate ports in a LAG.

STEP 1 Click Configuration > Port Management > Link Aggregation > LAGs. Information for each defined LAG is displayed.

STEP 2 Select the Load Balance Method:

- by MAC Address—(Default) Based on the destination and source MAC addresses of all packets.
- by IP and MAC Address—Based on the destination and source IP addresses for IP packets, and destination and source MAC addresses for non-IP packets.

STEP 3 Select the LAG to be configured, and click Edit.

STEP 4 Enter the values for the following fields:

- Operational Status—Displays the following:
- Status—Whether the LAG is currently operating.
- LAG Speed—Displays the current speed at which the LAG is operating.
- Flow Control—Whether flow control is enabled on the LAG.
- Port List—Move those ports that are to be assigned to the LAG from the Port List to the LAG Port Member list. Up to eight ports per static LAG can be assigned, and 16 ports can be assigned to a dynamic LAG.
- LAG Mode—Displays whether the LAG is up or down.
- Suspended LAG—Select to reactivate the LAG.
- LACP—Select to enable LACP on the selected LAG. This makes it a dynamic LAG. This field can only be enabled after moving a port to the LAG in the next field.
- Protected LAG—Select to make the LAG a protected port for Layer 2 isolation. See the Port Configuration description in Setting Basic Port Configuration for details regarding protected ports and LAGs.
- Auto Negotiation—Select to enable auto-negotiation on the LAG. Auto-negotiation is a protocol between two link partners that enables a LAG to advertise its transmission speed and flow control to its partner (the Flow Control default is disabled). It is recommended to keep auto-negotiation enabled on both sides of an aggregate link, or disabled on both sides, while ensuring that link speeds are identical.
- Port Speed—Configure the speed of the LAG. The port types determine the available speeds. You can designate this field only when port auto-negotiation is disabled.
- Auto Advertisement—Select the capabilities to be advertised by the LAG.

The options are as follows:

- Max Capability—All LAG speeds and both duplex modes are available.
- 10 Full Duplex—The LAG advertises a 10 Mbps speed and the mode is full duplex.
- 100 Full Duplex—The LAG advertises a 100 Mbps speed and the mode is full duplex.
- 1000 Full Duplex—The LAG advertises a 1000 Mbps speed and the mode is full duplex.
- Flow Control—Set Flow Control to either Enable or Disable or Auto-Negotiation.
- Description—Enter the LAG name or a comment.

STEP 5 Click Apply. LAG membership is saved to the Running Configuration file.

Green Ethernet

This section describes Green Ethernet, a set of features designed to be environmentally friendly by reducing the power consumption of a device.

The Green Ethernet feature can reduce overall power usage in the following ways:

- Energy-Detect Mode—In this mode, the switch conserves power when the operational status of a port is down. Energy-Detect Mode is supported on all ports.
- Short-Reach Mode—In the mode,, the switch will analyze cable length and adjust power usage accordingly. If the cable is shorter than 50 meters (164 feet), the device uses less power to send frames over the cable. This mode is only supported on RJ45 GE ports, and does not apply to Combo ports.

  This mode is globally disabled by default. It cannot be enabled if EEE mode is enabled (see below).

  802.3 Energy Efficient Ethernet (EEE)—EEE reduces power consumption when there is no traffic on the port. See Energy Efficient Ethernet Feature for more information.

  EEE is enabled globally by default. On a given port, if EEE is enabled, Short-Reach mode will be disabled. If Short Reach-Mode is enabled, EEE is grayed out.

  These modes are configured per port, without taking into account the LAG membership of the ports.

Power savings, current power consumption and cumulative energy saved can be monitored. The total amount of saved energy can be viewed as a percentage of the power that would have been consumed by the physical interfaces had they not been running in Green Ethernet mode.

The saved energy displayed does not include the amount of energy saved by EEE.

Energy Efficient Ethernet Feature

EEE is designed to save power when there is no traffic on the link. In Energy Detect Mode, power is reduced when the port is down.

When using 802.3 EEE, systems on both sides of the link can disable portions of their functionality and save power during periods of no traffic.

802.3 EEE supports IEEE 802.3 MAC operation at 100 Mbps and 1000 Mbps:

LLDP is used to select the optimal set of parameters for both devices. If LLDP is not supported by the link partner, or is disabled, 802.3 EEE will still be operational, but it might not be in the optimal operational mode.

The 802.3 EEE feature is implemented using a port mode called Low Power Idle (LPI) mode. The switch automatically chooses LPI Mode, if enabled, for a port when there is no traffic on that port.

Both sides of a connection (device port and connecting device) must support

802.3 EEE for it to work. When traffic is absent, both sides send signals indicating that power is about to be reduced. When signals from both sides are received, the Keep Alive signal indicates that the ports are in LPI Mode (and not in Down status), and power is reduced.

For ports to stay in LPI mode, the Keep Alive signal must be received continuously from both sides.

Power Saving by Disabling Port LEDs

The Disable Port LEDs feature saves power consumed by the device's LEDs. When located in an unoccupied room, these LEDs are unnecessary. Use the Green Ethernet feature to disable port LEDs (link, speed, and PoE) when they are not needed. Enable them if needed (debugging, connecting additional devices, etc.).

Advertise Capabilities Negotiation

802.3 EEE support is advertised during the Auto-Negotiation stage. Auto-Negotiation provides a linked device with the capability to detect the abilities (modes of operation) supported by the device at the other end of the link, determine common abilities, and configure itself for joint operation. Auto-Negotiation is performed at the time of link-up, on command from management, or upon detection of a link error. During the link establishment process, both link partners exchange their 802.3 EEE capabilities. Auto-Negotiation functions without user interaction when it is enabled on the device.

NOTE:
If Auto-Negotiation is not enabled on a port, the EEE is disabled. The only exception is if the link speed is 1GB, then EEE will still be enabled even though Auto-Negotiation is disabled.

Availability of 802.3 EEE

- Please check the release notes for a complete listing of products that support EEE.

Default Configuration

- By default, 802.3 EEE is enabled globally and per port.

Interactions Between Features

The following describe 802.3 EEE interactions with other features:

- If auto-negotiation is not enabled on the port, the 802.3 EEE operational status is disabled. The exception to this rule is that if the link speed is 1gigabyte, EEE will still be enabled even though Auto-Negotiation is disabled.
- If 802.3 EEE is enabled and the port is going up, it commences to work immediately in accordance with the maximum wake time value of the port.
- On the GUI, the EEE field for the port is not available when the Short Reach Mode option on the port is checked.
- If the port speed on the GE port is changed to 10Mbit, 802.3 EEE is disabled. This is supported in GE models only.

802.3 EEE Configuration Workflow

- This section describes how to configure the 802.3 EEE feature and view its counters.

STEP 1 Ensure that auto-negotiation is enabled on the port by opening the Ports page.

a. Select a port and open the Edit Ports page.

b. Select Auto Negotiation field to ensure that it is Enabled.

STEP 2 Ensure that 802.3 Energy Efficient Ethernet (EEE) is globally enabled in the Green Ethernet page (it is enabled by default). This page also displays how much energy has been saved.

STEP 3 Ensure that 802.3 EEE is enabled on a port by opening the Green Ethernet page.

a. Select a port, open the Edit Ports page.

b. Check the 802.3 Energy Efficient Ethernet (EEE) mode on the port (it is enabled by default).

Configuring Green Ethernet

To configure Green Ethernet globally and on a port, do the following:

STEP 1 Click Configuration > Port Management > Green Ethernet.

STEP 2 Choose whether to enable the following features:

- Energy Detect Mode—Select to globally enable.
- Port LEDs—Select to disable port LEDs. When disabled, ports do not display link status, activity, etc.
- Short Reach—Select to globally enable Short Reach mode if there are Green Ethernet ports on the device.

NOTE:
If Short Reach is enabled, EEE must be disabled.

- 802.3 Energy Efficient Ethernet (EEE)—Select to globally enable EEE.

STEP 3 Click Apply to set the global settings.

The following fields are displayed:

- Power Savings—Displays the percentage of power saved by running Port LED, Short Reach and Energy Detect modes. The EEE power savings is dynamic by nature since it is based on port utilization and is therefore not taken into consideration. The power saving calculation is performed by comparing the maximum power consumption without power savings to the current consumption.
- Cumulative Energy Saved—Displays the amount of energy saved from the last device reboot in watt hours. This value is updated each time there is an event that affects power saving.

For each port the following fields are described:

- Port—The port number.
- Short-Reach Mode—Whether Short-Reach Mode is enabled.
- Short Reach Status—Whether Short-Reach Mode is operational. This is a function of whether it has been enabled (Administrative Status), whether it has been enabled on the local port, and whether it is operational on the local port.
- Short Reach Reason—If Short-Reach mode is not operational, displays the reason.
- Cable Length—Displays VCT-returned cable length in meters.
- EEE Mode—Whether the mode is enabled.
- EEE Status—Whether EEE is currently operating on the local port. This is a function of whether it has been enabled (Administrative Status), whether it has been enabled on the local port and whether it is operational on the local port.

NOTE:
The window displays the Short Reach, Energy Detect and EEE settings for each port; however, they are not enabled on any port unless they are also enabled globally.

- EEE Status—State of EEE on the device (enabled or disabled).
- Remote EEE Mode—EEE mode of the linked partner.
- Energy Detect Mode—Whether Energy Detect Mode is enabled.
- Energy Detect Status—Whether Energy Detect Mode is currently operational. This is a function of whether it has been enabled (Administrative Status), whether it has been enabled on the local port and whether it is operational on the local port.

- Energy Detect Reason— If Energy Detect Mode is not operational, this field identifies why not.

STEP 5 Select a Port and click Edit.

STEP 6 Select to enable or disable the various features.

STEP 7 Click Apply. The Green Ethernet port settings are written to the Running Configuration File.

PoE

The Power over Ethernet (PoE) feature is only available on PoE-based devices. For a list of PoE-based devices, refer to the Device Models section.

This section describes how to use the PoE feature. It covers the following topics:

- Overview
- Feature Configuration
- Port Limit Power Mode
- Class Limit Power Mode

Overview

A PoE device is PSE (Power Sourcing Equipment) that delivers electrical power to connected PD (Powered Devices) over existing copper cables without interfering with the network traffic, updating the physical network or modifying the network infrastructure.

See Device Models for information concerning PoE support on various models. PoE provides the following features:

- Eliminates the need to run 110/220 V AC power to all devices on a wired LAN.
- Removes the necessity for placing all network devices next to power sources.
- Eliminates the need to deploy double cabling systems in an enterprise, significantly decreasing installation costs.

Power over Ethernet can be used in any enterprise network that deploys relatively low-powered devices connected to the Ethernet LAN:

- IP phones
- Wireless access points
- IP gateways
- Audio and video remote monitoring devices

PoE Operation

- PoE implementation stages are as follows:
- Detection—Sends special pulses on the copper cable. When a PoE device is located at the other end, that device responds to these pulses.
- Classification—Negotiation between the Power Sourcing Equipment PSE) and the Powered Device (PD) commences after the Detection stage. During negotiation, the PD specifies its class, which is the amount of maximum power that the PD consumes.
- Power Consumption—After the classification stage completes, the PSE provides power to the PD. If the PD supports PoE, but without classification, it is assumed to be class 0 (the maximum). If a PD tries to consume more power than permitted by the standard, the PSE stops supplying power to the port.

Power Modes

- Power per port can be limited depending on the Power Mode:
- Port Limit—Power is limited to a specified wattage. For these settings to be active, the system must be in PoE Port Limit mode. That mode is configured in the PoE Feature Configuration page. When the power consumed on the port exceeds the port limit, the port power is turned off.
- Class Limit—Power is limited based on the class of the connected PD. For these settings to be active, the system must be in PoE Class Limit mode. That mode is configured in the PoE Feature Configuration page. When the power consumed on the port exceeds the class limit, the port power is turned off.

PoE Priority Example

A 48-port device is supplying a total of 375 watts.

The administrator configures all ports to allocate up to 30 watts each. This results in 48 times 30 ports equaling 1440 watts, which is too much. The device cannot provide enough power to each port, so it provides power according to the priority.The administrator sets the priority for each port, allocating how much power it can be given.

These priorities are entered in the PoE Port Limit Mode or Class Limit Power Mode pages.

See Device Models for a description of the device models that support PoE and the maximum power that can be allocated to PoE ports.

PoE Configuration Considerations

- There are two factors to consider in PoE configuration:
- The amount of power that the PSE can supply
- The amount of power that the PD is attempting to consume

You can decide the following:

- Maximum power a PSE is allowed to supply to a PD
- POE mode—To change the mode from Class Power Limit to Port Limit, and vice versa, during device operation. The power values per port that were configured for the Port Limit mode are retained.

NOTE:
Changing the mode from Class Limit to Port limit, and vice versa, when the device is operational forces the Powered Device to reboot.

- Maximum port limit allowed as a per-port numerical limit in mW (Port Limit mode).
- The PoE-specific hardware automatically detects the PD class and its power limit according to the class of the device connected to each specific port (Class Limit mode).
- If at any time during the connectivity an attached PD requires more power from the device than the configured allocation allows (no matter if the device is in Class Limit or Port Limit mode), the devicedoes the following:
- Maintains the up/down status of the PoE port link
- Turns off power delivery to the PoE port
- Logs the reason for turning off power

CAUTION
Consider the following when connecting switches capable of supplying PoE:

The PoE model of the device is PSE (Power Sourcing Equipment) that is capable of supplying DC power to attaching PD (Powered Devices). These devices include VoIP phones, IP cameras, and wireless access points. The PoE switches can detect and supply power to pre-standard legacy PoE Powered Devices. Due to the support of legacy PoE, it is possible that a PoE device acting as a PSE may mistakenly detect and supply power to an attaching PSE, including other PoE switches, as a legacy PD.

Even though the PoE switches are PSE, and as such should be powered by AC, they could be powered up as a legacy PD by another PSE due to false detection. When this happens, the PoE device may not operate properly and may not be able to properly supply power to its attaching PDs.

To prevent false detection, you should disable PoE on the ports on the PoE switches that are used to connect to PSEs. You should also first power up a PSE device before connecting it to a PoE device. When a device is being falsely detected as a PD, you should disconnect the device from the PoE port and power cycle the device with AC power before reconnecting its PoE ports.

Feature Configuration

The Feature Configuration page enables selecting either the Port Limit or Class Limit PoE mode.

These settings are entered in advance. When the PD actually connects and is consuming power, it might consume much less than the maximum power allowed.

Output power is disabled during power-on reboot, initialization, and system configuration to ensure that PDs are not damaged.

To configure PoE on the device and monitor current power usage:

STEP 1 Click Configuration > Port Management > PoE > Feature Configuration.

STEP 2 Enter the values for the following fields:

- Power Mode—Select one of the following options:
- Port Limit—The maximum power limit per each port is configured by the user.
- Class Limit—The maximum power limit per port is determined by the class of the device, which results from the Classification stage.

NOTE:
When you change from Port Limit to Class Limit, or vice versa, you must disable PoE ports, and enable them after changing the power configuration.

The following counters are displayed for the device:

- Nominal Power —The total amount of power in watts that the device can supply to all the connected PDs.
- Consumed Power—Amount of power in watts that is currently being consumed by the PoE ports.
- Available Power—Nominal power in watts minus the amount of consumed power.

STEP 3 Click Apply to save the PoE properties.

Port Limit Power Mode

To configure port limit power mode do the following:

STEP 1 Click Configuration > Port Management > PoE > Port Limit Power Mode. The list of fields below is for Port Limit Power Mode.

The following fields are displayed for ports on which the port limit power mode is enabled:

- PoE Status—Enable or disable PoE on the port.
- Power Priority Level—Port priority is low, high, or critical, for use when the power supply is low. For example, if the power supply is running at 99% usage and port 1 is prioritized as high, but port 3 is prioritized as low, port 1 receives power and port 3 might be denied power.
- Power Allocation Limit (mW)—Power in milliwatts allocated to the port.
- Max Power Allocation (mW)—Maximum amount of power permitted on this port.
- Power Consumption (mW)—Amount of power assigned to the powered device connected to the selected interface.
- Class—Power class of device.
- Operational Status—Displays whether Power Limit mode is enabled or disabled on the port.

STEP 2 Select a port and click Edit. Enter the fields as described above.

STEP 3 Click Apply. The PoE settings for the port are written to the Running Configuration file.

Class Limit Power Mode

- To configure class limit power mode, do the following:

STEP 1 Click Configuration > Port Management > PoE > Class Limit Power Mode.

The following fields are displayed for ports on which the port limit power mode is enabled:

- PoE Status—Enable or disable PoE on the port.
- Power Priority Level—Port priority is low, high, or critical, for use when the power supply is low. For example, if the power supply is running at 99% usage and port 1 is prioritized as high, but port 3 is prioritized as low, port 1 receives power and port 3 might be denied power.
- Class—Class configured on this port. The classes are shown in the following:

• Class	• Maximum Power Delivered by Device Port
• 0	• 15.4 watt
• 1	• 4.0 watt
• 2	• 7.0 watt
• 3	• 15.4 watt
• 4	• 30.0 watt

- Max Power Allocation (mW)—Maximum amount of power permitted on this port.
- Power Consumption (mW)—Amount of power assigned to the powered device connected to the selected interface.
- Operational Status—Whether the Class Limit mode is enabled or disabled on the port.

STEP 2 Select a port and click Edit. Enter the fields as described above.

STEP 3 Click Apply. The PoE settings for the port are written to the Running Configuration file.

Chapter 6 VLAN Management

VLAN Management

This section covers the following topics:

- Overview
- VLANs
- Interfaces
- VLAN Memberships
- Voice VLAN

Overview

A VLAN is a logical group of ports that enables devices associated with it to communicate with each other over the Ethernet MAC layer, regardless of the physical LAN segment of the bridged network to which they are connected.

Each VLAN is configured with a unique VID (VLAN ID) with a value from 1 to 4094. A port on a device in a bridged network is a member of a VLAN if it can send data to and receive data from the VLAN. A port is an untagged member of a VLAN if all packets destined for that port into the VLAN have no VLAN tag. A port is a tagged member of a VLAN if all packets destined for that port into the VLAN have a VLAN tag. A port can be a member of one untagged VLAN and can be a member of several tagged VLANs.

A port in VLAN Access Mode can be part of only one VLAN. If it is in Trunk Mode, the port can be part of one or more VLANs.

VLANs address security and scalability issues. Traffic from a VLAN stays within the VLAN, and terminates at devices in the VLAN. It also eases network configuration by logically connecting devices without physically relocating those devices.

If a frame is VLAN-tagged, a four-byte VLAN tag is added to each Ethernet frame. The tag contains a VLAN ID between 1 and 4094, and a VLAN Priority Tag (VPT) between 0 and 7. See Quality of Service for details about VPT.

When a frame enters a VLAN-aware device, it is classified as belonging to a VLAN, based on the four-byte VLAN tag in the frame.

If there is no VLAN tag in the frame or the frame is priority-tagged only, the frame is classified to the VLAN based on the PVID (Port VLAN Identifier) configured at the ingress port where the frame is received.

Frames belonging to a VLAN remain within the VLAN. This is achieved by sending or forwarding a frame only to egress ports that are members of the target VLAN. An egress port may be a tagged or untagged member of a VLAN. The egress port peforms the following actions:

- Adds a VLAN tag to the frame if the egress port is a tagged member of the target VLAN, and the original frame does not have a VLAN tag.
- Removes the VLAN tag from the frame if the egress port is an untagged member of the target VLAN, and the original frame has a VLAN tag.

VLAN Roles

VLANs function at Layer 2. All VLAN traffic (Unicast/Broadcast/ Multicast) remains within its VLAN. Devices attached to different VLANs do not have direct connectivity to each other over the Ethernet MAC layer. Devices from different VLANs can communicate with each other only through Layer 3 routers. An IP router, for example, is required to route IP traffic between VLANs if each VLAN represents an IP subnet.

The IP router might be a traditional router, where each of its interfaces connects to only one VLAN. Traffic to and from a traditional IP router must be VLAN untagged. The IP router can be a VLAN-aware router, where each of its interfaces can connect to one or more VLANs. Traffic to and from a VLAN-aware IP router can be VLAN tagged or untagged.

Adjacent VLAN-aware devices exchange VLAN information with each other by using Generic VLAN Registration Protocol (GVRP). As a result, VLAN information is propagated through a bridged network.

VLANs on a device can be created statically or dynamically, based on the GVRP information exchanged by devices. A VLAN can be static or dynamic (from GVRP), but not both.

Some VLANs can have additional roles:

- Voice VLAN—For more information refer to Voice VLAN.
- Guest VLAN—For more information refer to Security: Network Access Control.
- Default VLAN—For more information refer to VLANs.
- Management VLAN (in Layer 2-system-mode systems) —For more information refer to the Layer 2 IP Addressing section.

VLAN Configuration Workflow

To configure VLANs, do the following:

1. If required, change the default VLAN by using the VLAN Management section.
2. Create the required VLANs by using the VLANs section.
3. Set the desired VLAN-related configuration for ports.
4. Assign interfaces to VLANs by using the Configuring VLAN Membership section.
5. View the current VLAN port membership for all the interfaces in the Configuring VLAN Membership section.

VLAN Creation

When using factory default settings, the device automatically creates VLAN 1 as the default VLAN, the default interface status of all ports is Trunk, and all ports are configured as untagged members of the default VLAN.

The default VLAN has the following characteristics:

- It is distinct, non-static/non-dynamic, and all ports are untagged members by default.
- It cannot be deleted.
- It cannot be given a label.
- It cannot be used for any special role, such as unauthenticated VLAN or Voice VLAN. This is only relevant for OUI-enabled voice VLAN.
- If a port is no longer a member of any VLAN, the device automatically configures the port as an untagged member of the default VLAN. A port is no longer a member of a VLAN if the VLAN is deleted or the port is removed from the VLAN.

When the VID of the default VLAN is changed, the device performs the following on all the ports in the VLAN, after saving the configuration and rebooting the device:

- Removes VLAN membership of the ports from the original default VLAN (possible only after reboot).
- Changes the PVID (Port VLAN Identifier) of the ports to the VID of the new default VLAN.
- The original default VLAN ID is removed from the device. To be used, it must be recreated.
- Adds the ports as untagged VLAN members of the new default VLAN.

VLANs

The VLANs page enables you to change the default VLAN and create a new VLAN.

To change the default VLAN, do the following:

STEP 1 Click Configuration > VLAN Management > VLANs.

STEP 2 Enter the value for the following field:

- Current Default VLAN—Displays the current default VLAN ID.
- Default VLAN After Reboot—Enter a new VLAN ID to replace the default VLAN after reboot.

STEP 3 Click Apply.

STEP 4 Click Save (in the upper-right corner of the window) and save the Running Configuration to the Startup Configuration.

The Default VLAN After Reboot becomes the Current Default VLAN after you reboot the device.

You can create a VLAN, but this has no effect until the VLAN is attached to at least one port, either manually or dynamically. Ports must always belong to one or more VLANs.

The device supports up to 1000 VLANs, including the default VLAN.

Each VLAN must be configured with a unique VID (VLAN ID) with a value from 1 to 4094. The device reserves VID 4095 as the Discard VLAN. All packets classified to the Discard VLAN are discarded at ingress, and are not forwarded to a port.

To create a VLAN, do the following:

STEP 1 Click Configuration > VLAN Management > VLANs.

This page displays the following fields for all VLANs:

- VLAN ID—User-defined VLAN ID.
- VLAN Name—User-defined VLAN name.
- Type—VLAN type
- Static—VLAN is user-defined.
- Default—VLAN is the default VLAN.

STEP 2 Click Add to add a new VLAN.

The page enables the creation of either a single VLAN or a range of VLANs.

STEP 3 Enter the following fields for the new VLAN(s).

- VLAN—Select one of the following options:
- Single VLAN—Select to create a single VLAN.
- Range of VLANs—Select to create a range of VLANs and specify the range of VLANs to be created by entering the Starting VID and Ending VID, inclusive. When using the Range function, the maximum number of VLANs you can create at one time is 100.
- VLAN ID—Enter a VLAN ID.
- VLAN Name—Enter a VLAN name.
- VLAN ID Range—Enter a range of VLANs.

STEP 4 Click Apply to create the VLAN(s).

Interfaces

The Interfaces page displays and enables configuration of VLAN-related parameters for all interfaces. To configure the VLAN settings:

STEP 1 Click Configuration > VLAN Management > Interfaces.

The following VLAN parameters are displayed for each interface:

- Interface VLAN Mode—The options are as follows:
- Access—The interface is an untagged member of a single VLAN. A port configured in this mode is known as an access port.
- Trunk—The interface is an untagged member of one VLAN at most, and is a tagged member of zero or more VLANs. A port configured in this mode is known as a trunk port.
- PVID—Port VLAN ID (PVID) of the VLAN to which incoming untagged and priority tagged frames are classified. The possible values are 1 to 4094.

STEP 2 To configure a port or LAG, select it and click Edit.

STEP 3 Enter the values for the Interface VLAN Mode and PVID

STEP 4 Click Apply. The parameters are written to the Running Configuration file.

VLAN Memberships

The VLAN Memberships page displays the VLAN memberships of the ports in various presentations. You can use them to add memberships to or remove memberships from the VLANs.

When a port is forbidden default VLAN membership, that port is not allowed membership in any other VLAN. An internal VID of 4095 is assigned to the port.

To forward the packets properly, intermediate VLAN-aware devices that carry VLAN traffic along the path between end nodes must either be manually configured or must dynamically learn the VLANs and their port memberships from Generic VLAN Registration Protocol (GVRP).

Untagged port membership between two VLAN-aware devices with no intervening VLAN-aware devices, must be to the same VLAN. In other words, the PVID on the ports between the two devices must be the same if the ports are to send and receive untagged packets to and from the VLAN. Otherwise, traffic might leak from one VLAN to another.

Frames that are VLAN-tagged can pass through other network devices that are VLAN-aware or VLAN-unaware. If a destination end node is VLAN-unaware, but is to receive traffic from a VLAN, then the last VLAN-aware device (if there is one), must send frames of the destination VLAN to the end node untagged.

Use the VLAN Memberships page to display and configure the ports within a specific VLAN.

To map ports or LAGs to a VLAN, do the following:

STEP 1 Click Configuration > VLAN Management > VLAN Memberships.

STEP 2 Select a VLAN ID and the Interface Type (Port or LAG), and click Search to display or to change the port characteristic with respect to the VLAN.

The port mode for each port or LAG appears with its current port mode (Access or Trunk) configured from the Interfaces page.

Each port or LAG appears with its current registration to the VLAN.

STEP 3 Change the registration of an interface to the VLAN by selecting the desired option from the following list:

- PVID—Select to set the PVID of the interface to the VID of the VLAN. PVID is a per-port setting.
- Access—Select to make the interface an access interface on this VLAN.
- Trunk—Select to make the interface a trunk interface on this VLAN.
- Forbidden—The interface is not allowed to join the VLAN even from GVRP registration. When a port is not a member of any other VLAN, enabling this option on the port makes the port part of internal VLAN 4095 (a reserved VID).
- Excluded—The interface is currently not a member of the VLAN. This is the default for all the ports and LAGs. The port can join the VLAN through GVRP registration.
- Tagged—The interface is a tagged member of the VLAN.
- Untagged—The interface is an untagged member of the VLAN. Frames of the VLAN are sent untagged to the interface VLAN.

STEP 4 Click Apply. The interfaces are assigned to the VLAN, and written to the Running Configuration file.

You can continue to display and/or configure port membership of another VLAN by selecting another VLAN ID.

Voice VLAN

This section covers the following topics:

- Overview
- Feature Configuration
- Telephony OUI Interfaces

Overview

In a LAN, voice devices, such as IP phones, VoIP endpoints, and voice systems are placed into the same VLAN. This VLAN is referred as the voice VLAN.
If the voice devices are in different voice VLANs, IP (Layer 3) routers are needed to provide communication.

The device supports a single voice VLAN. By default, the voice VLAN is VLAN 1. A different voice VLAN can be manually configured.

Telephony Organizationally Unique Identifier

Telephony OUIs are assigned by the Institute of Electrical and Electronics Engineers, Incorporated (IEEE) Registration Authority. Since the number of IP phone manufacturers is limited and well-known, the known OUI values cause the relevant frames, and the port on which they are seen, to be automatically assigned to the Voice VLAN. The OUI Global table can hold up to 128 OUIs.

When a port is manually configured as a candidate to join the voice VLAN, the device dynamically adds the port to the voice VLAN if it receives a packet with a source MAC address matching to one of the configured telephony OUIs. An OUI is the first three bytes of an Ethernet MAC address.

Voice End-Points

For a voice VLAN work properly, voice devices, such as IP phones and VoIP endpoints, must be assigned to the voice VLAN where it sends and receives its voice traffic. Some of the possible scenarios are as follows:

- A phone/endpoint may be statically configured with the voice VLAN.
- A phone/endpoint may obtain the voice VLAN in the boot file it downloads from a TFTP server. A DHCP server may specify the boot file and the TFTP server when it assigns an IP address to the phone.

The device expects the attaching voice devices to send voice VLAN, tagged packets. On ports where the voice VLAN is also the native VLAN, voice VLAN untagged packets are possible.

Voice VLAN QoS

Working with the OUI mode, the device can additionally configure the mapping and remarking (CoS/802.1p) of the voice traffic based on the OUI.

By default, all interfaces are CoS/802.1p trusted. The device applies the quality of service based on the CoS/802.1p value found in the voice stream. For Telephony OUI voice streams, you can override the quality of service and optionally remark the 802.1p of the voice streams by specifying the desired CoS/802.1p values and using the remarking option under Telephony OUI.

Voice VLAN Constraints

The following constraints exist:

- Only one Voice VLAN is supported.
- A VLAN that is defined as a Voice VLAN cannot be removed
- The Voice VLAN cannot be the default VLAN.
- The Voice VLAN cannot support DVA (Dynamic VLAN assignment).
- The Voice VLAN cannot be the Guest VLAN if the voice VLAN mode is OUI. If the voice VLAN mode is Auto, then the Voice VLAN can be the Guest VLAN.
- The Voice VLAN QoS decision has priority over any other QoS decision, except for the Policy/ACL QoS decision.
- A new VLAN ID can be configured for the Voice VLAN only if the current Voice VLAN does not have candidate ports.
- The interface VLAN of a candidate port must be in Trunk mode.
- The Voice VLAN QoS is applied to candidate ports that have joined the Voice VLAN, and to static ports.
- The voice flow is accepted if the MAC address can be learned by the Forwarding Database (FDB). (If there is no free space in FDB, no action occurs).

Voice VLAN Workflow

Workflow1: To configure the Telephony OUI Method

STEP 1 Configure Telephony OUI in the Feature Configuration page.

STEP 2 Configure Telephony OUI VLAN membership for ports in the Telephony OUI page.

Feature Configuration

To configure Voice VLAN and add OUIs to the Telephone OUI table, do the following:

STEP 1 Click Configuration > VLAN Management > Voice VLAN > Feature Configuration.

STEP 2 Enter values for the following fields:

- Voice VLAN ID—Select the VLAN that is to be the Voice VLAN.
- CoS/802.1p —Select a CoS/802.1p value .

Telephone OUI

- Telephone OUI Voice VLAN —Select to enable voice VLAN by OUI.
- Remark CoS/802 1p—Select to remark the CoS or 802.1p value.
- Aging Time—Enter the time delay to remove a port from the voice VLAN after all of the MAC addresses of the phones detected on the ports have aged out.

STEP 3 Click Apply. The VLAN properties are written to the Running Configuration file.

STEP 4 To add a new OUI, click Add.

STEP 5 Enter the values for the following fields:

- Telephony OUI—Enter a new OUI.
- Description—Enter the OUI name.

Telephony OUI Interfaces

The QoS attributes can be assigned per port to the voice packets in one of the following modes:

- All—Quality of Service (QoS) values configured to the Voice VLAN are applied to all of the incoming frames that are received on the interface and are classified to the Voice VLAN.
- Telephony Source MAC Address (SRC)—The QoS values configured for the Voice VLAN are applied to any incoming frame that is classified to the Voice VLAN and contains an OUI in the source MAC address that matches a configured telephony OUI.

Use the Telephony OUI page to add an interface to the voice VLAN on the basis of the OUI identifier and to configure the OUI QoS mode of voice VLAN.

To configure Telephony OUI on an interface, do the following:

STEP 1 Click Configuration > VLAN Management > Voice VLAN > Telephony OUI Interfaces.

The Telephony OUI Interfaces page contains voice VLAN OUI parameters for all interfaces.

STEP 2 To configure an interface to be a candidate port of the telephony OUI-based voice VLAN, click Edit.

STEP 3 Enter the values for the following fields:

- Interface—Select an interface.
- Telephony OUI —Select to indicate that the interface is a candidate port of the telephony OUI based voice VLAN. When packets that match one of the configured telephony OUI are received, the port is added to the voice VLAN.
- QoS Mode—Select one of the following options:

- All—QoS attributes are applied on all packets that are classified to the Voice VLAN. *Telephony Source MAC Address*-QoS attributes are applied only on packets from IP phones.
- *Telephony Source MAC Address*-QoS attributes are applied only on packets from IP phones.

STEP 4 Click **Apply**. The OUI is added.

Chapter 7 Spanning Tree Management

Spanning Tree Management

This section describes the Spanning Tree Protocol (STP) (IEEE802.1D and IEEE802.1Q) and covers the following topics:

- Overview
- Spanning Tree
- STP Interfaces
- RSTP Interfaces

Overview

STP protects a Layer 2 broadcast domain from broadcast storms by selectively setting links to standby mode to prevent loops. In standby mode, these links temporarily stop transferring user data. After the topology changes so that the data transfer is made possible, the links are automatically re-activated.

Loops occur when alternate routes exist between hosts. Loops in an extended network can cause switches to forward traffic indefinitely, resulting in increased traffic load and reduced network efficiency.

STP provides a tree topology for any arrangement of switches and interconnecting links, by creating a unique path between end stations on a network, and thereby eliminating loops.

The device supports the following Spanning Tree Protocol versions:

- Classic STP – Provides a single path between any two end stations, avoiding and eliminating loops.
- Rapid STP (RSTP) – Detects network topologies to provide faster convergence of the spanning tree. This is most effective when the network topology is naturally tree-structured, and therefore faster convergence might be possible. RSTP is enabled by default.

Spanning Tree

The STP Status and Global Settings page contains parameters for enabling STP or RSTP.

Use the STP Interface Settings page, RSTP Interface Settings page to configure each mode, respectively.

To set the STP status and global settings, do the following:

STEP 1 Click Configuration > Spanning Tree Management > Spanning Tree.

STEP 2 Enter the parameters.

Global Settings

- Spanning Tree—Select to enable STP on the device.
- Spanning Tree Mode—Select an STP mode—Classic STP or Rapid STP.
- Path Cost Default Values—Selects the method used to assign default path costs to the STP ports. The default path cost assigned to an interface varies according to the selected method.
- Short—Specifies the range 1 through 65,535 for port path costs.
- Long—Specifies the range 1 through 200,000,000 for port path costs.
- BPDU Handling—Select how Bridge Protocol Data Unit (BPDU) packets are managed when STP is disabled on the port or the device. BPDUs are used to transmit spanning tree information.
- Filtering—Filters BPDU packets when Spanning Tree is disabled on an interface.
- Flooding—Floods BPDU packets when Spanning Tree is disabled on an interface.

Bridge Configuration

- Priority—Set the global priority value. The priority value influences the port choice when a bridge has two ports connected in a loop. The priority is a value from 0 to 240, set in increments of 16.

- Hello Time—Set the interval (in seconds) that a root bridge waits between configuration messages.
- Maximum Age—Set the interval (in seconds) that the device can wait without receiving a configuration message, before attempting to redefine its own configuration.
- Forward Delay—Set the interval (in seconds) that a bridge remains in a learning state before forwarding packets.

Status Designated Root

- Bridge ID—The combined bridge priority and the MAC address of the device.
- Root Bridge ID—The combined root bridge priority and the MAC address of the root bridge.
- Root Port—The port that offers the lowest cost path from this bridge to the root bridge. (This is significant when the bridge is not the root.)
- Root Path Cost—The cost of the path from this bridge to the root.
- Topology Changes Count—The total number of STP topology changes that have occurred.
- Last Topology Change—The time interval that elapsed since the last topology change occurred. The time appears in a days/hours/minutes/seconds format.

STEP 3 Click Apply. The STP Global settings are written to the Running Configuration file.

STP Interfaces

The STP Interface page enables you to configure STP on a per-port basis, and to view the information learned by the protocol, such as the designated bridge.

The defined configuration entered is valid for all flavors of the STP protocol. To configure STP on an interface, do the following:

STEP 1 Click Configuration > Spanning Tree Management > STP Interfaces.

STEP 2 Select an interface type and click Edit.

STEP 3 Enter the parameters.

- STP—Select to enable STP on the port.
- BPDU Handling—Select how BPDU packets are managed when STP is disabled on the port or the device. BPDUs are used to transmit spanning tree information.
- Use Global Settings—Select to use the settings defined in the Spanning Tree page.
- Filtering—Filters BPDU packets when Spanning Tree is disabled on an interface.
- Flooding—Floods BPDU packets when Spanning Tree is disabled on an interface.
- Path Cost—Set the port contribution to the root path cost or use the default cost generated by the system.
- Priority—Select the priority value of the port. The priority value influences the port choice when a bridge has two ports connected in a loop. The priority is a value from 0 to 240, set in increments of 16.
- Port State—Displays the current STP state of a port.
- Disabled—STP is currently disabled on the port. The port forwards traffic while learning MAC addresses.
- Blocking—The port is currently blocked, and cannot forward traffic (with the exception of BPDU data) or learn MAC addresses.
- Listening—The port is in Listening Mode. The port cannot forward traffic, and cannot learn MAC addresses.
- Learning—The port is in Learning Mode. The port cannot forward traffic, but it can learn new MAC addresses.

- Forwarding—The port is in Forwarding Mode. The port can forward traffic and learn new MAC addresses.
- Designated Bridge ID—Displays the priority and interface of the selected port.
- Designated Port ID—Displays the priority and interface of the selected port.
- Designated Cost—Displays the cost of the port participating in the STP topology. Ports with a lower cost are less likely to be blocked if STP detects loops.
- Root Guard—Enables or disables Root Guard on the device. The Root Guard option provides a way to enforce the root bridge placement in the network.

Root Guard ensures that the port on which this feature is enabled is the designated port. Normally, all root bridge ports are designated ports, unless two or more ports of the root bridge are connected. If the bridge receives superior BPDUs on a Root Guard-enabled port, Root Guard moves this port to a root-inconsistent STP state. This root-inconsistent state is effectively equal to a listening state. No traffic is forwarded across this port. In this way, Root Guard enforces the position of the root bridge.

- BPDU Guard—Enables or disables the Bridge Protocol Data Unit (BPDU) Guard feature on the port.

The BPDU Guard enables you to enforce the STP domain borders and keep the active topology predictable. The devices behind the ports that have BPDU Guard enabled cannot influence the STP topology. At the reception of BPDUs, the BPDU guard operation disables the port that has BPDU configured. In this case, a BPDU message is received, and an appropriate SNMP trap is generated.

STEP 4 Click Apply. The interface settings are written to the Running Configuration file.

RSTP Interfaces

Rapid Spanning Tree Protocol (RSTP) enables a faster STP convergence without creating forwarding loops.

The RSTP Interfaces page enables you to configure RSTP per port. Any configuration that is done on this page is active when the global STP mode is set to RSTP

To enter RSTP settings, do the following:

STEP 1 Click Configuration > Spanning Tree Management> RSTP Interfaces.

STEP 2 Select a port.

STEP 3 Select an interface, and click Edit.

STEP 4 Enter the Interface Settings.

- Point to Point Mode — Define the point-to-point link status. Ports defined as full duplex are considered point-to-point port links.
- Enable — This port is an RSTP edge port when this feature is enabled, and is brought to Forwarding Mode quickly (usually within 2 seconds).
- Disable — The port is not considered point-to-point for RSTP purposes, which means that STP works on it at regular speed, as opposed to high speed.
- Auto — Automatically determines the device status by using RSTP BPDUs.
- Edge Port Mode — Enables or disables Fast Link on the port. If Fast Link Mode is enabled on a port, the port is automatically set to forwarding state when the port link is up. Fast Link optimizes the STP protocol convergence. The options are:
- Enable — Enables Fast Link immediately.
- Disable — Disables Fast Link.
- Auto — Enables Fast Link a few seconds after the interface becomes active. This allows STP to resolve loops before enabling Fast Link.

NOTE:
It is recommended to set the value to Auto so that the device sets the port to fast link mode if a host is connected to it, or sets it as a regular STP port if connected to another device. This helps avoid loops.

- STP Mode — Select either STP or RSTP.
- Point to Point Status—Displays the point-to-point operational status if the Point to Point Administrative Status is set to Auto.
- Port Role — Displays the role of the port that was assigned by STP to provide STP paths. The possible roles are as follows:
- Root — Lowest cost path to forward packets to the root bridge.
- Designated — The interface through which the bridge is connected to the LAN, which provides the lowest cost path from the LAN to the root bridge.
- Alternate — Provides an alternate path to the root bridge from the root interface.
- Backup — Provides a backup path to the designated port path toward the spanning tree leaves. This provides a configuration in which two ports are connected in a loop by a point-to-point link. Backup ports are also used when a LAN has two or more established connections to a shared segment.
- Disabled — The port is not participating in spanning tree.
- Port Status — Displays the RSTP status on the specific port.
- Disabled — STP is currently disabled on the port.
- Blocking — The port is currently blocked, and it cannot forward traffic or learn MAC addresses.
- Listening — The port is in Listening Mode. The port cannot forward traffic, and cannot learn MAC addresses.
- Learning — The port is in Learning Mode. The port cannot forward traffic, however it can learn new MAC addresses.
- Forwarding — The port is in Forwarding Mode. The port can forward traffic and learn new MAC addresses.

STEP 5 Click Apply. The Running Configuration file is updated.

Chapter 8 MAC Address Management

MAC Address Management

This section describes how to add MAC addresses to the system. It covers the following topics:

- Overview
- Dynamic MAC Addresses

Overview

A new source MAC address that appears in a frame arriving at the device is added to the Dynamic Address table. This MAC address is retained for a configurable period of time. If another frame with the same source MAC address does not arrive at the device before that time period expires, the MAC entry is aged (deleted) from the table.

When a frame arrives at the device, the device searches for a corresponding/matching destination MAC address entry in the dynamic table. If a match is found, the frame is marked for egress on the port specified in the table. If frames are sent to a MAC address that is not found in the tables, they are transmitted/broadcasted to all the ports on the relevant VLAN. Such frames are referred to as unknown Unicast frames.

The device supports a maximum of 8K dynamic MAC addresses.

Dynamic MAC Addresses

The Dynamic Address Table (bridging table) contains the MAC addresses acquired by monitoring the source addresses of frames entering the device.

To prevent this table from overflowing, and to make room for new MAC addresses, an address is deleted if no corresponding traffic is received for a certain period. This period of time is the aging interval.

Configuring Dynamic MAC Address Aging Time

To configure the aging interval for dynamic addresses, do the following:

STEP 1 Click Configuration > MAC Address Management > Dynamic MAC Addresses.

STEP 2 Enter Aging Time. The aging time is a value between the user-configured value and twice that value minus 1. For example, if you entered 300 seconds, the aging time is between 300 and 599 seconds.

STEP 3 Click Apply. The aging time is updated.To query dynamic addresses, do the following:

STEP 4 In the Dynamic MAC Address Table block, enter the query criteria:

- VLAN ID—Enter the VLAN ID for which the table is queried.
- MAC Address—Enter the MAC address for which the table is queried.
- Interface—Select the interface for which the table is queried. The query can search for specific unit/slot, ports, or LAGs.
- Sort By—Select the field for which the table is queried.

STEP 5 Click Search. The Dynamic MAC Address Table is queried and the results are displayed.To delete all dynamic MAC addresses click Clear.

Chapter 9 Multicast

Multicast

This section describes the Multicast Forwarding feature and covers the following topics:

- Overview
- Feature Configuration
- IGMP Snooping
- Multicast Router Ports
- Forward All
- Unregistered Multicast
- IGMP IP Group Addresses
- MAC Group Address FDB
- IP Group Address FDB

Overview

Multicast forwarding enables one-to-many information dissemination. Multicast applications are useful for dissemination of information to multiple clients, where clients do not require reception of the entire content. A typical application is a cable-TV-like service, where clients can join a channel in the middle of a transmission, and leave before it ends.

The data is sent only to relevant ports. Forwarding the data only to the relevant ports conserves bandwidth and host resources on links.

For Multicast forwarding to work across IP subnets, nodes and routers must be Multicast-capable. A Multicast-capable node must be able to do the following:

- Send and receive Multicast packets.
- Register the Multicast addresses being listened to by the node with local routers, so that local and remote routers can route the Multicast packet to the nodes.

Typical Multicast Setup

While Multicast routers route Multicast packets between IP subnets, Multicast-capable Layer 2 switches forward Multicast packets to registered nodes within a LAN or VLAN.

A typical setup involves a router that forwards the Multicast streams between private and/or public IP networks, a device with Internet Group Membership Protocol (IGMP) snooping capabilities, and a Multicast client that wants to receive a Multicast stream. In this setup, the router sends IGMP queries periodically.

These queries reach the device, which in turn floods the queries to the VLAN, and also learns the port where there is a Multicast router (Mrouter). When a host receives the IGMP query message, it responds with an IGMP Join message saying that the host wants to receive a specific Multicast stream and optionally from a specific source. The device with IGMP snooping analyzes the Join messages, and learns that the Multicast stream the host has requested must be forwarded to this specific port. It then forwards the IGMP Join to the Mrouter only. Similarly, when the Mrouter receives an IGMP Join message, it learns the interface from which it received the Join messages that wants to receive a specific Multicast stream. The Mrouter forwards the requested Multicast stream to the interface.

In a Layer 2 Multicast service, a Layer 2 switch receives a single frame addressed to a specific Multicast address. It creates copies of the frame to be transmitted on each relevant port.

When the device is IGMP snooping-enabled and receives a frame for a Multicast stream, it forwards the Multicast frame to all the ports that have registered to receive the Multicast stream using IGMP Join messages.

The device can forward Multicast streams based on one of the following options:

- Multicast MAC Group Address
- IP Multicast Group Address (G)
- A combination of the source IP address (S) and the destination IP Multicast Group Address (G) of the Multicast packet.
- One of these options can be configured per VLAN.

The system maintains lists of Multicast groups for each VLAN, and this manages the Multicast information that each port should receive. The Multicast groups and their receiving ports can be configured statically or learned dynamically using IGMP snooping.

Multicast registration is the process of listening and responding to Multicast registration protocols. The available protocols are IGMP for IPv4.

When IGMP snooping is enabled in a device on a VLAN, it analyzes the IGMP packets it receives from the VLAN connected to the device and Multicast routers in the network.

When a device learns that a host is using IGMP messages to register to receive a Multicast stream, optionally from a specific source, the device adds the registration to its Multicast Forwarding Data Base (MFDB).

IGMP snooping can effectively reduce Multicast traffic from streaming bandwidth- intensive IP applications. A device using IGMP snooping only forwards Multicast traffic to the hosts interested in that traffic. This reduction of Multicast traffic reduces the packet processing at the device, and also reduces the workload of the end hosts, since they do not have to receive and filter all of the Multicast traffic generated in the network.

The following versions are supported: IGMP v1/v2/v3.

Multicast Address Properties

Multicast addresses have the following properties:

- Each IPv4 Multicast address is in the address range 224.0.0.0 to 239.255.255.255.
- To map an IP Multicast group address to a Layer 2 Multicast address: for IPv4, this is mapped by taking the 23 low-order bits from the IPv4 address, and adding them to the 01:00:5e prefix. By standard, the upper nine bits of the IP address are ignored, and any IP addresses that only differ in the value of these upper bits are mapped to the same Layer 2 address, since the lower 23 bits that are used are identical. For example, 234.129.2.3 is mapped to a MAC Multicast group address 01:00:5e:01:02:03. Up to 32 IP Multicast group addresses can be mapped to the same Layer 2 address.

Feature Configuration

The Feature Configuration page enables you to configure the Bridge Multicast filtering status.

By default, all Multicast frames are flooded to all ports of the VLAN. To selectively forward only to relevant ports and filter (drop) the Multicast on the rest of the ports, enable Bridge Multicast filtering status in the Feature Configuration page.

If filtering is enabled, Multicast frames are forwarded to a subset of the ports in the relevant VLAN as defined in the Multicast Forwarding Data Base. Multicast filtering is enforced on all traffic. By default, such traffic is flooded to all relevant ports, but you can limit forwarding to a smaller subset.

A common way of representing Multicast membership is the (S,G) notation where S is the (single) source sending a Multicast stream of data, and G is the IPv4 group address. If a Multicast client can receive Multicast traffic from any source of a specific Multicast group, this is saved as (*,G).

The following are ways of forwarding Multicast frames:

- MAC Group Address—Based on the destination MAC address in the Ethernet frame.

NOTE:
One or more IP Multicast group addresses can be mapped to a MAC group address. Forwarding, based on the MAC group address, can result in an IP Multicast stream being forwarded to ports that have no receiver for the stream.

- IP Group Address—Based on the destination IP address of the IP packet (*,G).
- Source Specific IP Group Address—Based on both the destination IP address and the source IP address of the IP packet (S,G).

By selecting the forwarding mode, you can define the method used by hardware to identify Multicast flow by one of the following options: MAC Group Address, IP Group Address, or Source Specific IP Group Address. (S,G) is supported by IGMPv3, while IGMPv1/2 support only (*.G), which is just the group ID. The device supports a maximum of 256 static and dynamic Multicast group addresses. To enable Multicast filtering, and select the forwarding method:

STEP 1 Click Configuration > Multicast > Feature Configuration.

STEP 2 Enter the global parameter:

- Bridge Multicast Filtering—Select to enable filtering of Multicast addresses.

VLAN Settings

- VLAN ID—Select the VLAN ID to set its forwarding method.
- IPv4 Multicast Forwarding—Select one of the following options:
- By MAC Address—Select to enable the MAC address method for forwarding Multicast packets.
- By IPv4 Group Address—Select to enable the IPv4 group address method for forwarding Multicast packets.
- By Source Specific IPv4 Group Address—Select to enable the source-specific IPv4 group address method for forwarding Multicast packets.

STEP 3 Click Apply. The Running Configuration file is updated.

IGMP Snooping

To support selective Multicast forwarding (IPv4), Bridge Multicast filtering must be enabled (in the Feature Configuration page), and IGMP Snooping must be enabled globally and for each relevant VLAN (in the IGMP Snooping page).

By default, a Layer 2 device forwards Multicast frames to all ports of the relevant VLAN, essentially treating the frame as if it were a Broadcast. With IGMP Snooping the device forwards Multicast frames to ports that have registered Multicast clients.

NOTE:
The device supports IGMP Snooping only on static VLANs. It does not support IGMP Snooping on dynamic VLANs.

When IGMP Snooping is enabled globally or on a VLAN, all IGMP packets are forwarded to the CPU. The CPU analyzes the incoming packets, and determines the following:

- Which ports are asking to join which Multicast groups on what VLAN.
- Which ports are connected to Multicast routers (Mrouters) that are generating IGMP queries.
- Which ports are receiving PIM, DVMRP, or IGMP query protocols. These are displayed on the IGMP Snooping page.

Ports, asking to join a specific Multicast group, issue an IGMP report that specifies which group(s) the host wants to join. This results in the creation of a forwarding entry in the Multicast Forwarding Data Base.

To enable IGMP Snooping and identify the device as an IGMP Snooping Querier on a VLAN, do the following:

STEP 1 Click Configuration > Multicast > IGMP Snooping.

STEP 2 Enable or disable IGMP Snooping.

When IGMP Snooping is enabled globally, the device monitoring network traffic can determine which hosts have requested to receive Multicast traffic.

The device only performs IGMP Snooping if both IGMP snooping and Bridge Multicast filtering are enabled.

STEP 3 Select a VLAN, and click Edit. Enter the parameters:

- VLAN ID—Select the VLAN ID on which IGMP snooping is defined.

VLAN Settings

- IGMP Snooping Status—Enable or disable the monitoring of network traffic for the selected VLAN.
- Auto Learn MRouter Ports —Select to enable auto learning of the ports to which the Mrouter is connected.
- Immediate Leave—Select to enable Immediate Leave to decrease the time it takes to block a Multicast stream sent to a member port when an IGMP Group Leave message is received on that port.
- IGMP Querier—Select to enable the IGMP Querier.
- IGMP Querier Version—Select the IGMP version used if the device becomes the elected querier. Select IGMPv3 if there are switches and/or Multicast routers in the VLAN that perform source-specific IP Multicast forwarding.
- Querier Source IP Address—Select the source IP address of the IGMP Querier. The following options are available:
- Auto—The system decides whether to use the IP address of the VLAN or the management IP address.
- User Defined—This can be the IP address of the VLAN or it can be the management IP address.

STEP 4 Click Apply. The Running Configuration file is updated.

Multicast Router Ports

A Multicast router (Mrouter) port is a port that connects to a Multicast router. The device includes the Multicast router port(s) numbers when it forwards the Multicast streams and IGMP registration messages. This is required so that the Multicast routers can, in turn, forward the Multicast streams and propagate the registration messages to other subnets.

To statically configure or see dynamically-detected ports connected to the Multicast router, do the following:

STEP 1 Click Configuration > Multicast > Multicast Router Ports.

STEP 2 Enter some or all of following query filter criteria:

- VLAN ID—Select the VLAN ID for the router ports that are described.
- Interface Type—Select whether to display ports or LAGs.

STEP 3 Click Search. The interfaces matching the query criteria are displayed.

STEP 4 For each port or LAG, select its association type. The options are as follows:

- Static—The port is statically configured as a Multicast router port.
- Dynamic—(Display only) The port is dynamically configured as a Multicast router port by a IGMP query. To enable the dynamic learning of Multicast router ports, go to the IGMP Snooping page.
- Forbidden—This port is not to be configured as a Multicast router port, even if IGMP queries are received on this port. If Forbidden is enabled on a port, Mrouter is not learned on this port (i.e. MRouter Ports Auto-Learn is not enabled on this port).
- None—The port is not currently a Multicast router port.

STEP 5 Click Apply to update the device.

Forward All

The Forward All page enables and displays the configuration of the ports and/or LAGs that are to receive Multicast streams from a specific VLAN. This feature requires that Bridge Multicast filtering in the Feature Configuration page be enabled. If it is disabled, then all Multicast traffic is flooded to ports in the device.

You can statically (manually) configure a port to Forward All, if the devices connecting to the port do not support IGMP.

IGMP messages are not forwarded to ports defined as Forward All.

NOTE:
The configuration affects only the ports that are members of the selected VLAN.

To define Forward All Multicast, do the following:

STEP 1 Click Configuration > Multicast > Forward All.

STEP 2 Define the following:

- VLAN ID — The VLAN ID the ports/LAGs are to be displayed.
- Interface Type — Define whether to display ports or LAGs.

STEP 3 Click Search. The status of all ports/LAGs are displayed.

STEP 4 Select the port/LAG that is to be defined as Forward All by using the following methods:

- Static — The port receives all Multicast streams.
- Forbidden — Ports cannot receive any Multicast streams, even if IGMP
- snooping designated the port to join a Multicast group.
- None — The port is not currently a Forward All port.

STEP 5 Click Apply. The Running Configuration file is updated.

Unregistered Multicast

Multicast frames are generally forwarded to all ports in the VLAN. If IGMP Snooping is enabled, the device learns about the existence of Multicast groups, and monitors which ports have joined which Multicast group. Multicast groups can also be statically configured. Multicast groups that were either dynamically learned or statically configured are considered registered.

The device forwards Multicast frames (from a registered Multicast group) only to ports that are registered to that Multicast group.

The Unregistered Multicast page enables handling Multicast frames that belong to groups that are not known to the device (unregistered Multicast groups). Unregistered Multicast frames are usually forwarded to all ports on the VLAN.

You can select a port to receive or filter unregistered Multicast streams. The configuration is valid for any VLAN of which it is a member (or will be a member).

This feature ensures that the customer receives only the Multicast groups requested and not others that may be transmitted in the network.

To define unregistered Multicast settings, do the following:

STEP 1 Click Configuration > Multicast > Unregistered Multicast.

STEP 2 Define the following:

- Interface Type — Define whether to display ports or LAGs.
- Interface Settings — Displays the forwarding status of the selected interface. The possible values are as follows:
- Forwarding — Enables forwarding of unregistered Multicast frames to the selected interface.
- Filtering — Enables filtering (rejecting) of unregistered Multicast frames to the selected interface.

STEP 3 Click Apply. The settings are saved, and the Running Configuration file is updated.

IGMP IP Group Addresses

The IGMP IP Group Addresses page displays the IPv4 group address learned from IGMP messages.

There might be a difference between information on this page and, for example, information displayed in the MAC Group Address FDB page. Assuming that the system is in MAC-based groups and a port that requested to join the following Multicast groups 224.1.1.1 and 225.1.1.1, both are mapped to the same MAC Multicast address 01:00:5e:01:01:01. In this case, there is a single entry in the MAC Group Address FDB page, but two entries on this page.

To query for an IP Multicast group, do the following:

STEP 1 Click Configuration > Multicast > IGMP IP Group Addresses.

STEP 2 Enter some or all of following query filter criteria:

- VLAN ID — Defines the VLAN ID to query.
- IP Group Address — Defines the Multicast group MAC address or IP address to query.
- Source IP Address — Defines the sender address to query.

STEP 3 Click Search. The following fields are displayed for each Multicast group:

- VLAN ID — The VLAN ID.
- IP Group Address — The Multicast group MAC address or IP address.
- Source IP Address — The sender address for all of the specified group ports.
- Included Ports — The list of destination ports for the Multicast stream.
- Excluded Ports — The list of ports not included in the group.
- Compatibility Mode — The oldest IGMP version of registration from the hosts the device receives on the IP group address.

MAC Group Address FDB

The device supports forwarding incoming Multicast traffic based on the Multicast group information. This information is derived from the IGMP packets received or as the result of manual configuration, and it is stored in the Multicast Forwarding Database (MFDB).

When a frame is received from a VLAN that is configured to forward Multicast streams, based on MAC group addresses, and its destination address is a Layer 2 Multicast address, the frame is forwarded to all ports that are members of the MAC group address.

The MAC Group Address FDB page has the following functions:

- Query and view information from the MFDB, relating to a specific VLAN ID or a specific MAC address group. This data is acquired either dynamically through IGMP snooping or statically by manual entry.
- Add or delete static entries to the MFDB that provide static forwarding information, based on MAC destination addresses.
- Display a list of all ports/LAGs that are a member of each VLAN ID and MAC address group, and enter whether traffic is forwarded to it or not.

To define and view MAC Multicast groups, do the following:

STEP 1 Click Configuration > Multicast > MAC Group Address FDB.

STEP 2 Enter the parameters.

- VLAN ID —Enter the VLAN ID of the group to be displayed.
- MAC Group Address —Set the MAC address of the Multicast group to be displayed. If no MAC Group Address is specified, the page contains all the MAC Group Addresses from the selected VLAN.

STEP 3 Click Search, and the MAC Multicast group addresses are displayed in the lower block.

Entries that were created both in this page and in the IP Group Address FDB page are displayed. For those created in the IP Group Address FDB page, the IP addresses are converted to MAC addresses.

STEP 4 Click Add to add a static MAC Group Address.

STEP 5 Enter the parameters.

- VLAN ID — Defines the VLAN ID of the new Multicast group.
- MAC Group Address — Defines the MAC address of the new Multicast group.

STEP 6 Click Apply, the MAC Multicast group is saved to the Running Configuration file.

To configure and display the registration for the interfaces within the group, select an address, and click Membership.

The MAC Group Address FDB page opens. Enter the following:

- VLAN ID—The VLAN ID of the Multicast group.
- MAC Group Address—The MAC address of the group.
- Interface Type—Port or LAG.

STEP 7 Click Search to display the port or LAG membership.

STEP 8 Select the way that each interface is associated with the Multicast group:

- Static—Attaches the interface to the Multicast group as a static member.
- Dynamic—Indicates that the interface was added to the Multicast group as a result of IGMP snooping.
- Forbidden—Specifies that this port is not allowed to join this group on this VLAN.
- Excluded—Specifies that the port is not currently a member of this Multicast group on this VLAN.

STEP 9 Click Apply, and the Running Configuration file is updated.

NOTE:
Entries that were created in the IP Group Address FDB page cannot be deleted in this page (even if they are selected).

IP Group Address FDB

The IP Group Address FDB page enables querying and adding IP Multicast groups contained in the IP Multicast Groups Forwarding Data Base.

To define and view IP Multicast groups, do the following:

STEP 1 Click Configuration > Multicast > IP Group Address FDB

The page contains all of the IP Multicast group addresses learned by snooping.

STEP 2 Enter the parameters required for filtering.

- VLAN ID—Enter the VLAN ID of the group to be displayed.
- IP Group Address—Define the IP address of the Multicast group to be displayed. This is only relevant when the Forwarding Mode is (S,G).
- Source IP Address—Define the source IP address of the sending device. If mode is (S,G), enter the sender S. This together with the IP group address is the Multicast group ID (S,G) to be displayed. If mode is (*.G), enter an * to indicate that the Multicast group is only defined by destination.

STEP 3 Click Search. The results are displayed in the lower block.

STEP 4 Click Add to add a static IP Multicast group address.

STEP 5 Enter the parameters.

- VLAN ID—Defines the VLAN ID of the group to be added.
- IP Group Address—Define the IP address of the new Multicast group.

Group Address Settings

- Source Specific IP Multicast—Select to indicate that the entry contains a specific source, and adds the address in the IP Source Address field. If not, the entry is added as a (*,G) entry, an IP group address from any IP source.
- Source IP Address—Enter the source address to be included.

STEP 6 Click Apply. The IP Multicast group is added, and the device is updated.

STEP 7 To configure and display the registration of an IP group address, select an address and click Membership.

The VLAN ID, IP Version, IP Multicast group address, and Source IP address selected are displayed as read-only in the top of the window. You can select the filter type:

- Interface Type—Select whether to display ports or LAGs.

STEP 8 For each interface, select its association type. The options are as follows:

- Static—Attaches the interface to the Multicast group as a static member.
- Dynamic—Indicates that the interface was added to the Multicast group as a result of IGMP snooping.
- Forbidden—Specifies that this port is forbidden from joining this group on this VLAN.
- Excluded—Indicates that the port is not currently a member of this Multicast group on this VLAN. This is selected by default until Static or Forbidden is selected.

STEP 9 Click Apply. The Running Configuration file is updated.

Chapter 10 IP Interface

IP Interface

This section describes IP interfaces and covers the following topics:

- IPv4
- IPv6
- ARP

IPv4

This section describes IPv4 configuration. It covers the following topics:

- Overview
- IPv4 Interface<managed>in Layer 2 System Mode
- IPv4 Static Routes<managed>

Overview

Some features are only available in Layer 2 or Layer 3 system mode, as described below:

- In Layer 2 system mode, the device operates as a Layer 2 VLAN-aware device, and has no routing capabilities.
- In Layer 3 system mode, the device has IP routing capabilities and Layer 2 system mode capabilities. In this system mode, a Layer 3 port still retains much of the Layer 2 functionality, such as Spanning Tree Protocol and VLAN membership.

In Layer 3 system mode, the device does not support MAC-based VLAN, Dynamic VLAN Assignment, VLAN Rate Limit, SYN Rate DoS Protection, or Advanced QoS Policers.

Configuring the device to work in either mode is performed in the Maintenance > System Mode & Reboot page.

NOTE:
Switching from one system mode (layer) to another requires a mandatory reboot, and the startup configuration of the device is then deleted.

Layer 2 IP Addressing

In Layer 2 system mode, the device has one IPv4 address and up to two IPv6 interfaces (either "native" interface or Tunnel) in the management VLAN. This IP address and the default gateway can be configured manually, or by DHCP. The static IP address and default gateway for Layer 2 system mode are configured on the IPv4 Interface and IPv6 Interfaces pages. In Layer 2 system mode, the device uses the default gateway, if configured, to communicate with devices that are not in the same IP subnet with the device. By default, VLAN 1 is the management VLAN, but this can be modified. When operating in Layer 2 system mode, the device can only be reached at the configured IP address through its management VLAN.

The factory default setting of the IPv4 address configuration is DHCPv4. This means that the device acts as a DHCPv4 client, and sends out a DHCPv4 request during boot up.

If the device receives a DHCPv4 response from the DHCPv4 server with an IPv4 address, it sends Address Resolution Protocol (ARP) packets to confirm that the IP address is unique. If the ARP response shows that the IPv4 address is in use, the device sends a DHCPDECLINE message to the offering DHCP server, and sends another DHCPDISCOVER packet that restarts the process.

If the device does not receive a DHCPv4 response in 60 seconds, it continues to send DHCPDISCOVER queries, and adopts the default IPv4 address: 192.168.1.251/24.

IP address collisions occur when the same IP address is used in the same IP subnet by more than one device. Address collisions require administrative actions on the DHCP server and/or the devices that collide with the device.

When a VLAN is configured to use dynamic IPv4 addresses, the device issues DHCPv4 requests until it is assigned an IPv4 address from a DHCPv4 server. In Layer 2 system mode, only the management VLAN can be configured with a static or dynamic IP address. In Layer 3 system mode, all the interface types (ports, LAGs, and/or VLANs) on the device can be configured with a static or dynamic IP address.

The IP address assignment rules for the device are as follows:

- When in Layer 2 system mode, unless the device is configured with a static IP address, it issues DHCPv4 requests until a response is received from the DHCP server.
- If the IP address on the device is changed, the device issues gratuitous ARP packets to the corresponding VLAN to check IP address collisions. This rule also applies when the device reverts to the default IP address.
- The system status LED changes to solid blue when a new unique IP address is received from the DHCP server. If a static IP address has been set, the system status LED also changes to solid blue. The LED flashes when the device is acquiring an IP address and is currently using the factory default IP address 192.168.1. 251.
- The same rules apply when a client must renew the lease, prior to its expiration date through a DHCPREQUEST message.
- With factory default settings, when no statically defined or DHCP-acquired IP address is available, the default IP address is used. When the other IP addresses become available, the addresses are automatically used. The default IP address is always on the management VLAN.

Layer 3 IP Addressing

In Layer 3 system mode, the device can have multiple IP addresses. Each IP address can be assigned to specified ports, LAGs, or VLANs. These IP addresses are configured in the IPv4 Interface and IPv6 Interfaces pages in Layer 3 system mode. This provides more network flexibility than the Layer 2 system mode, in which only a single IP address can be configured. Operating in Layer 3 system mode, the device can be reached at all of its IP addresses from the corresponding interfaces.

A predefined, default route is not provided in Layer 3 system mode. To remotely manage the device, a default route must be defined. All DHCP-assigned default gateways are stored as default routes. In addition, you can manually define default routes. This is defined in the IPv4 Static Routes pages.

All the IP addresses configured or assigned to the device are referred to as

Management IP addresses in this guide.

If the pages for Layer 2 and Layer 3 are different, both versions are displayed.

IPv4 Interface

IPv4 interfaces can be defined on the device when it is in Layer 2 or

Layer 3 system mode.

IPv4 Interface in Layer 2 System Mode

To manage the device by using the web-based configuration utility, the IPv4 device management IP address must be defined and known. The device IP address can be manually configured or automatically received from a DHCP server.

To configure the IPv4 device IP address, do the following:

STEP 1 Click Configuration > IP Interface > IPv4 > IPv4 Interface.

STEP 2 Enter values for the following fields:

- Management VLAN—Select the Management VLAN used to access the device through telnet or the Web GUI. VLAN1 is the default Management VLAN.
- IP Address Type—Select one of the following options:
- Dynamic (DHCP)—Discover the IP address using DHCP from the management VLAN.
- Static IP Address—Manually define a static IP address.

NOTE:
DHCP Option 12 (Host Name option) is supported when the device is a DHCP client. If DHCP Option 12 is received from a DHCP server, it is saved as the server's host name. DHCP option 12 will not be requested by the device. The DHCP server must be configured to send option 12, regardless of what is requested in order to make use of this feature.

- Dynamic IP Address—Select to renew the DHCP-supplied IP address.
- IP Address—Enter the IP address, and configure one of the following Mask fields:
- IP Subnet Mask—Configure one of the following Mask fields:
- SubNet Mask—Select and enter the IP address mask.
- Prefix Length—Select and enter the length of the IPv4 address prefix.

- User Defined Default Gateway—Select User Defined and enter the default gateway IP address, or select None to remove the selected default gateway IP address from the interface.
- Default Gateway—Displays the current default gateway status.

NOTE:
If the device is not configured with a default gateway, it cannot communicate with other devices that are not in the same IP subnet.

STEP 3 Click Apply. The IPv4 interface settings are written to the Running Configuration file.

Defining IPv4 Interface in Layer 3 System Mode

The IPv4 Interface page is used when the device is in Layer 3 system mode. This mode enables configuring multiple IP addresses for device management, and provides routing services.

The IP address can be configured on a port, a LAG, or VLAN.

Operating in Layer 3 mode, the device routes traffic between the directly attached IP subnets configured on the device. The device continues to bridge traffic between devices in the same VLAN. Additional IPv4 routes for routing to non-directly attached subnets can be configured in the IPv4 Static Routes page.

NOTE:
The device software consumes one VLAN ID (VID) for every IP address configured on a port or LAG. The device takes the first VID that is not used starting from 4094.

To configure the IPv4 addresses, do the following:

STEP 1 Click Configuration > IP Interface > IPv4 > IPv4 Interface.

STEP 2 Select IPv4 Routing to enable the device to function as an IPv4 router.

STEP 3 Click Apply. The parameter is saved to the Running Configuration file.

This page displays the following fields in the IPv4 Interface Table:

- Interface—Interface for which the IP address is defined.
- IP Address—Configured IP address for the interface.
- IP Subnet Mask—Configured IP address mask.
- IP Address Type—IP address defined as static or DHCP.
- Dynamic IP Address—Received from DHCP server.
- Static—Entered manually.
- Status—Results of the IP address duplication check.
- Tentative—There is no final result for the IP address duplication check.
- Valid—The IP address collision check was completed, and no IP address collision was detected.
- Valid-Duplicated—The IP address duplication check was completed, and a duplicate IP address was detected.
- Duplicated—A duplicated IP address was detected for the default IP address.
- Delayed—The assignment of the IP address is delayed for 60 seconds if DHCP Client is enabled on startup to give time to discover DHCP address.
- Not Received—Relevant for DHCP Address. When a DHCP Client starts a discovery process, it assigns a dummy IP address 0.0.0.0 before the real address is obtained. This dummy address has the status of "Not Received"

STEP 4 Click Add. Enter the fields as described above.

STEP 5 Click Apply. The IPv4 address settings are written to the Running Configuration file.

IPv4 Static Routes

When the device is in Layer 3 system mode this page enables configuring and viewing IPv4 static routes on the device. When routing traffic, the next hop is decided according to the longest prefix match (LPM algorithm). A destination IPv4 address may match multiple routes in the IPv4 Static Route Table. The device uses the matched route with the highest subnet mask, that is, the longest prefix match.

To define an IP static route, do the following:

STEP 1 Click Configuration > Multicast > IPv4 > IPv4 Static Routes.

STEP 2 Click Add.

STEP 3 Enter values for the following fields:

- IP Subnet Address—Enter the destination IP address prefix.
- IP Subnet Mask—Select and enter information for one of the following:
- Network Mask—The IP route prefix for the destination IP.
- Prefix Length—The IP route prefix for the destination IP.
- Route Type—Select the route type.

- Reject—Rejects the route and stops routing to the destination network via all gateways. This ensures that if a frame arrives with the destination IP of this route, it is dropped.
- Remote—Indicates that the route is a remote path.
- Local—A directly connected network whose prefix is derived from a manually configured device's IPv6 address.
- Next Hop Router IP Address—Enter the next hop IP address or IP alias on the route.

NOTE:
You cannot configure a static route through a directly connected IP subnet where the device gets its IP address from a DHCP server.

- Metric—Enter the administrative distance to the next hop. The range is 1–255.

STEP 4 Click Apply. The IP Static route is saved to the Running Configuration file.

ARP

The device maintains an ARP (Address Resolution Protocol) table for all known devices that reside in the IP subnets directly connected to it. A directly-connected IP subnet is the subnet to which an IPv4 interface of the device is connected. When the device is required to send/route a packet to a local device, it searches the ARP table to obtain the MAC address of the device. The ARP table contains dynamic addresses. The device creates dynamic addresses from the ARP packets it receives. Dynamic addresses age out after a configured time.

NOTE:
The IP/MAC address mapping in the ARP Table is used to forward traffic originated by the device.

To define the ARP tables, do the following:

STEP 1 Click Configuration > IP Interface > IPv4 > ARP.

STEP 2 Enter the parameters.

- Interface—Select the interface for which to display information.
- ARP Entry Aging Time (1-40000000)—Enter the number of seconds that dynamic addresses can remain in the ARP table. A dynamic address ages out after the time it is in the table exceeds the ARP Entry Age Out time. When a dynamic address ages out, it is deleted from the table, and only returns when it is relearned.

STEP 3 Click Apply. The ARP global settings are written to the Running Configuration file.

The ARP table displays the following fields:

- IP Interface—The IPv4 Interface of the directly-connected IP subnet where the IP device resides.
- IP Address—The IP address of the IP device.
- MAC Address—The MAC address of the IP device.
- Status—Whether the entry was manually entered (static) or dynamically learned.

STEP 4 Click Add.

STEP 5 Enter the parameters:

Interface—An IPv4 interface can be configured on a port, LAG or VLAN. Select the desired interface from the list of configured IPv4 interfaces on the device.

- IP Address—Enter the IP address of the local device.
- MAC Address—Enter the MAC address of the local device.

STEP 6 Click Apply. The ARP entry is saved to the Running Configuration file.

IPv6

This section describes IPv6 configuration. It covers the following topics:

- Overview
- IPv6 Interfaces
- IPv6 Interface Addresses
- IPv6 Default Routers

Overview

The Internet Protocol version 6 (IPv6) is a network-layer protocol for packet-switched Internet works. IPv6 was designed to replace IPv4, the predominantly deployed Internet protocol.

IPv6 introduces greater flexibility in assigning IP addresses because the address size increases from 32-bit to 128-bit addresses. IPv6 addresses are written as eight groups of four hexadecimal digits, for example FE80:0000:0000:0000:0000:9C00 :876A:130B. The abbreviated form, in which a group of zeroes can be left out, and replaced with '::', is also acceptable, for example, ::-FE80::9C00:876A:130B.

IPv6 nodes require an intermediary mapping mechanism to communicate with other IPv6 nodes over an IPv4-only network. This mechanism, called a tunnel, enables IPv6-only hosts to reach IPv4 services, and enables isolated IPv6 hosts and networks to reach an IPv6 node over the IPv4 infrastructure.

The device detects IPv6 frames by the IPv6 EtherType.

IPv6 Interface

An IPv6 interface can be configured on a port, LAG, or VLAN. To define an IPv6 interface, do the following:

STEP 1 Click Configuration > IP Interface> IPv6 > IPv6 Interface.

STEP 2 Click Add to add a new interface on which interface IPv6 is enabled.

STEP 3 Enter the fields:

- IPv6 Interface—Select a specific port, LAG, or VLAN for the IPv6 address.

Interface Settings

- Number of DAD Attempts—Enter the number of consecutive neighbor solicitation messages that are sent while Duplicate Address Detection (DAD) is performed on the interface's Unicast IPv6 addresses. DAD verifies the uniqueness of a new Unicast IPv6 address before it is assigned. New addresses remain in a tentative state during DAD verification. Entering 0 in this field disables duplicate address detection processing on the specified interface. Entering 1 in this field indicates a single transmission without follow-up transmissions.
- IPv6 Address Auto Configuration—Select to enable automatic address configuration from router advertisements sent by neighbors.

NOTE:
The device does not support stateful address autoconfiguration from a DHCPv6 server.

- Send ICMPv6 Messages—Select to enable generating unreachable destination messages.

STEP 4 Click Apply to enable IPv6 processing on the selected interface. Regular IPv6 interfaces have the following addresses automatically configured:

- Link local address using EUI-64 format interface ID based on a device's MAC address
- All link local Multicast addresses (FF02::1)
- Solicited-Node Multicast address (format FF02::1:FFXX:XXXX)

IPv6 Interface Addresses

To assign an IPv6 address to an IPv6 Interface, do the following:

STEP 1 Click Configuration > IP Interface> IPv6 > IPv6 Interface Addresses.

STEP 2 To filter the table, select an interface name, and click Search. The interface appears in the IPv6 Address Table.

STEP 3 Click Add.

STEP 4 Enter values for the fields.

- IPv6 Interface—Displays the interface on which the IPv6 address is to be defined. If an * is displayed, the IPv6 interface is not enabled but has been configured.
- IPv6 Address Type—Select the type of the IPv6 address to add.
- Link Local—An IPv6 address that uniquely identifies hosts on a single network link. A link local address has a prefix of FE80, is not routable, and can be used for communication only on the local network. Only one link local address is supported. If a link local address exists on the interface, this entry replaces the address in the configuration.
- Global—An IPv6 address that is a global Unicast IPV6 type that is visible and reachable from other networks.
- IPv6 Address—In Layer 2, the device supports a singleIPv6 interface. In addition to the default link local and Multicast addresses, the device also automatically adds global addresses to the interface based on the router advertisements it receives. The device supports a maximum of 128 addresses at the interface. Each address must be a valid IPv6 address that is specified in hexadecimal format by using 16-bit values separated by colons.
- Prefix Length—The length of the Global IPv6 prefix is a value from 0-128 indicating the number of the high-order contiguous bits of the address that comprise the prefix (the network portion of the address).
- EUI-64—Select to use the EUI-64 parameter to identify the interface ID portion of the Global IPv6 address on a device MAC address.

STEP 5 Click Apply. The Running Configuration file is updated.

IPv6 Default Routers

The IPv6 Default Routers page enables configuring and viewing the default IPv6 router addresses. This list contains the routers that are candidates to become the device default router for nonlocal traffic (it may be empty). The device randomly selects a router from the list. The device supports one static IPv6 default router. Dynamic default routers are routers that have sent router advertisements to the device IPv6 interface.

When adding or deleting IP addresses, the following events occur:

- When removing an IP interface, all the default router IP addresses are removed. Dynamic IP addresses cannot be removed.
- An alert message appears after an attempt is made to insert more than a single user-defined address.
- An alert message appears when attempting to insert a non-link local type address, meaning 'fe80:'.

To define a default router, do the following:

STEP 1 Click Configuration > IP Interface> IPv6 > Default Routers.

This page displays the following fields for each default router:

- Default Router IPv6 Address — Link local IP address of the default router.
- IPv6 Interface — Outgoing IPv6 interface where the default router resides.
- State — Whether route is reachable or unreachable.
- Type — The default router configuration that includes the following options:
- Static — The default router was manually added to this table through the Add button.
- Dynamic — The default router was dynamically configured.

STEP 2 Click Add to add a static default router.

STEP 3 Enter the following fields:

- IPv6 Interface — Displays the outgoing Link Local interface.
- Default Router IPv6 Address — The IP address of the default router

STEP 4 Click Apply. The default router is saved to the Running Configuration file.

Chapter 11 IP Network Operations

IP Network Operations

This section covers the following topics:

- Domain Name System
- DHCP
- IP Source Guard<managed>
- ARP Inspection

Domain Name System

The Domain Name System (DNS) translates domain names into IP addresses for the purpose of locating and addressing hosts.

As a DNS client, the device resolves domain names to IP addresses through the use of one or more configured DNS servers.

DNS

Use the DNS page to enable the DNS feature, configure the DNS servers and set the default domain used by the device.

STEP 1 Click Configuration > IP Network Operations > Domain Name System > DNS.

STEP 2 Enter the following fields:

- DNS—Select to designate the device as a DNS client, which can resolve
- DNS names into IP addresses through one or more configured DNS servers.
- Default Domain Name—Enter the DNS domain name used to complete unqualified host names. The device appends this to all non-fully qualified domain names (NFQDNs) turning them into FQDNs.

The following fields are displayed for each configured DNS server:

- DNS Server IP Address—IP address of the DNS server.
- DNS Server State — Whether DNS server is Active or Inactive.
- IP Interface — Interface connected to DNS server.
- Preference — Each server has a preference value, a lower value means a higher chance of being used.
- Configuration Source — Source of the server's IP address (static or DHCPv4 or DHCPv6)

STEP 3 Up to eight DNS servers can be defined. To add a DNS server, click Add. Enter the parameters.

- IP Version—Select Version 6 for IPv6 or Version 4 for IPv4.
- IPv6 Address Type—Select the IPv6 address type (if IPv6 is used). The options are the following:
- Global — The IPv6 address is a global Unicast IPV6 type that is visible and reachable from other networks.
- Link Local — The IPv6 address uniquely identifies hosts on a single network link. A link local address has a prefix of FE80, is not routable, and can be used for communication only on the local network. Only one link local address is supported. If a link local address exists on the interface, this entry replaces the address in the configuration.
- Interface — If the IPv6 address type is Link Local, select the interface through which it is received.
- DNS Server IP Address — Enter the DNS server IP address.

Server Settings

- Preference — Select a value that determines the order in which the domains are used (from low to high). This effectively determines the order in which unqualified names are completed during DNS queries.

STEP 4 Click Apply. The DNS server is saved to the Running Configuration file.

DHCP

This section describes DHCP configuration. It covers the following topics:

- Overview
- DHCP Snooping/Relay
- DHCP Interfaces
- DHCP Snooping Trusted Interface

Overview

DHCP snooping provides a security mechanism to prevent receiving false DHCP response packets and to log DHCP addresses. It does this by treating ports on the device as either trusted or untrusted.

A trusted port is a port that is connected to a DHCP server and is allowed to assign DHCP addresses. DHCP messages received on trusted ports are allowed to pass through the device.

An untrusted port is a port that is not allowed to assign DHCP addresses. By default, all ports are considered untrusted until you declare them trusted (in the DHCP Snooping Trusted Interface page).

DHCPv4 Relay

DHCP Relay relays DHCP packets to the DHCP server.

DHCPv4 in Layer 2 and Layer 3

In Layer 2 system mode, the device relays DHCP messages received from VLANs on which DHCP Relay has been enabled.

In Layer 3 system mode, the device can also relay DHCP messages received from VLANs that do not have IP addresses. Whenever DHCP Relay is enabled on a VLAN without an IP address, Option 82 is inserted automatically. This insertion is in the specific VLAN, and does not influence the global administration state of Option 82 insertion.

Transparent DHCP Relay

For Transparent DHCP Relay where an external DHCP relay agent is being used, do the following:

- Enable DHCP Snooping.
- Enable Option 82 insertion.
- Disable DHCP Relay.

For regular DHCP Relay, do the following:

- Enable DHCP Relay.
- No need to enable Option 82 insertion.

Option 82

Option 82 (DHCP Relay Agent Information Option) passes port and agent information to a central DHCP server, indicating where an assigned IP address physically connects to the network.

The main goal of option 82 is to help to the DHCP server select the best IP subnet

(network pool) from which to obtain an IP address.

The following Option 82 options are available on the device:

- DHCP Insertion - Add Option 82 information to packets that do not have foreign Option 82 information.
- DHCP Passthrough - Forward or reject DHCP packets that contain Option 82 information from untrusted ports. On trusted ports, DHCP packets containing Option 82 information are always forwarded.

Interactions Between DHCPv4 Snooping, DHCPv4 Relay and Option 82

The following tables describe how the device behaves with various combinations of DHCP Snooping, DHCP Relay and Option 82.

The following describes how DHCP request packets are handled when DHCP Snooping is not enabled and DHCP Relay is enabled.

	DHCP Relay VLAN with IP Address		**DHCP Relay VLAN without IP Address**	
	Packet arrives without Option 82	Packet arrives with Option 82	Packet arrives without Option 82	Packet arrives with Option 82
Option 82 Insertion Disabled	Packet is sent without Option 82	Packet is sent with original Option 82	Relay - Inserts Option 82 Bridge – No Option 82 is inserted	Relay – Discards Option 82 Bridge – Packet is set with original Option 82
Option 82 Insertion Enabled	Relay – Is sent with Option 82 Bridge – No option 82 is sent	Packet is sent with original Option 82	Relay – Is sent with Option 82 Bridge – No option 82 is sent	Relay – Discards Option 82 Bridge – Packet is set with original Option 82

The following describes how DHCP request packets are handled when both DHCP snooping and DHCP relay are enabled.

	DHCP Relay VLAN with IP Address		**DHCP Relay VLAN without IP Address**	
	Packet arrives without Option 82	Packet arrives with Option 82	Packet arrives without Option 82	Packet arrives with Option 82
Option 82 Insertion Disabled	Packet is sent without Option 82	Packet is sent with original Option 82	Relay - Inserts Option 82 Bridge – No Option 82 is inserted	Relay – Discards Option 82 Bridge – Packet is sent with original Option 82
Option 82 Insertion Enabled	Relay – Is sent with Option 82 Bridge – No option 82 is added *If port is trusted, behaves as if DHCP Snooping is not enabled.*	Packet is sent with original Option 82	Relay – Is sent with Option 82 Bridge – Option 82 is added *If port is trusted, behaves as if DHCP Snooping is not enabled.*	Relay – Discards Option 82 Bridge – Packet is sent with original Option 82

The following describes how DHCP request packets are handled when DHCP snooping is disabled.

	DHCP Relay VLAN with IP Address		DHCP Relay VLAN without IP Address	
	Packet arrives without Option 82	Packet arrives with Option 82	Packet arrives without Option 82	Packet arrives with Option 82
Option 82 Insertion Disabled	Packet is sent without Option 82	Packet is sent with original Option 82	Relay - Discards Option 82 Bridge – Packet is sent without Option 82	If reply originates in the device, the packet is sent without Option 82. If reply does not originate in the device, the packet is discarded. Bridge – Packet is sent with original Option 82
Option 82 Insertion Enabled	Packet is sent without Option 82	Relay - Packet is sent without Option 82 Bridge – Packet is sent with Option 82	Relay – Discards Option 82 Bridge – Packet is sent without Option 82	Relay - Packet is sent without Option 82 Bridge – Packet is sent with Option 82

DHCP Snooping Binding Database

DHCP Snooping builds a database (known as the DHCP Snooping Binding database) derived from information taken from DHCP packets entering the device through trusted ports.

The DHCP Snooping Binding database contains the following data: input port, input VLAN, MAC address of the client, and IP address of the client if it exists

The DHCP Snooping Binding database is also used by IP Source Guard and

Dynamic ARP Inspection features to determine legitimate packet sources.

DHCP Trusted Ports

Ports can be either DHCP trusted or untrusted. By default, all ports are untrusted. To create a port as trusted, use the DHCP Snooping Trusted Interface page. Packets from these ports are automatically forwarded. Packets from trusted ports are used to create the Binding database and are handled as described below.

If DHCP Snooping is not enabled, all ports are trusted by default.

How the DHCP Snooping Binding Database is Built

The following describes how the device handles DHCP packets when both the DHCP client and DHCP server are trusted. The DHCP Snooping Binding database is built in this process.

DHCP Trusted Packet Handling

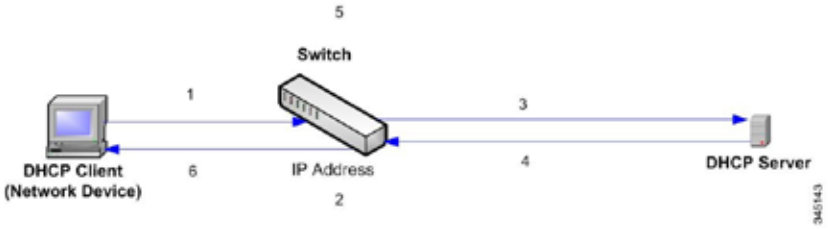

The actions are as follows:

STEP 1 Device sends DHCPDISCOVER to request an IP address or DHCPREQUEST to accept an IP address and lease.

STEP 2 Device snoops packet and adds the IP-MAC information to the DHCP Snooping Binding database.

STEP 3 Device forwards DHCPDISCOVER or DHCPREQUEST packets.

STEP 4 DHCP server sends DHCPOFFER packet to offer an IP address, DHCPACK to assign one, or DHCPNAK to deny the address request.

STEP 5 Device snoops packet. If an entry exists in the DHCP Snooping Binding table that matches the packet, the device replaces it with IP-MAC binding on receipt of DHCPACK.

STEP 6 Device forwards DHCPOFFER, DHCPACK, or DHCPNAK.

The following summarizes how DHCP packets are handled from both trusted and untrusted ports. The DHCP Snooping Binding database is stored in non-volatile memory.

DHCP Snooping Packet Handling

Packet Type	**Arriving from Untrusted Ingress Interface**	**Arriving from Trusted Ingress Interface**
DHCPDISCOVER	Forward to trusted interfaces only.	Forwarded to trusted interfaces only.
DHCPOFFER	Filter.	Forward the packet according to DHCP information. If the destination address is unknown the packet is filtered.
DHCPREQUEST	Forward to trusted interfaces only.	Forward to trusted interfaces only.
DHCPACK	Filter.	Same as DHCPOFFER and an entry is added to the DHCP Snooping Binding database.
DHCPNAK	Filter.	Same as DHCPOFFER. Remove entry if exists.
DHCPDECLINE	Check if there is information in the database. If the information exists and does not match the interface on which the message was received, the packet is filtered. Otherwise, the packet is forwarded to trusted interfaces only, and the entry is removed from database.	Forward to trusted interfaces only
Packet Type	Arriving from Untrusted Ingress Interface	Arriving from Trusted Ingress Interface
DHCPRELEASE	Same as DHCPDECLINE.	Same as DHCPDECLINE.
DHCPINFORM	Forward to trusted interfaces only.	Forward to trusted interfaces only.
DHCPLEASEQUERY	Filtered.	Forward.

DHCP Snooping Along With DHCP Relay

If both DHCP Snooping and DHCP Relay are globally enabled, then if DHCP Snooping is enabled on the client's VLAN, DHCP Snooping rules contained in the DHCP Snooping Binding database are applied. The DHCP Snooping Binding database is updated in the client's and DHCP server's VLAN for packets that are relayed.

DHCP Default Configuration

The following describes DHCP Snooping and DHCP Relay default options.

DHCP Default Options

Option	Default State
DHCP Snooping	Enabled
Option 82 Insertion	Not enabled
Option 82 Passthrough	Not enabled
Verify MAC Address	Enabled
Backup DHCP Snooping Binding Database	Not enabled
DHCP Relay	Disabled

Configuring DHCP Work Flow

To configure DHCP Relay and DHCP Snooping, do the following:

STEP 1 Enable DHCP Snooping and/or DHCP Relay in the Configuration > IP Network Operations > DHCP Relay & Snooping page.

STEP 2 Define the interfaces on which DHCP Snooping is enabled in the Configuration > IP Network Operations > DHCP>DHCP Interfaces page.

STEP 3 Configure interfaces as trusted or untrusted in the Configuration > IP Network Operations> DHCP > Trust Interfaces page.

STEP 4 Optional. Add entries to the DHCP Snooping Binding database in the Configuration >IP Network Operations> DHCP Snooping Binding Database page.

DHCP Snooping /Relay

This section describes how the DHCP Snooping and Relay features are implemented via the Web-based interface.

In Layer 2, DHCP Relay and Snooping can only be enabled on VLANs with IP addresses.

In Layer 3, DHCP Relay and Snooping can be enabled on any interface with an IP address, and on VLANs with or without an IP address.

To globally configure DHCP Snooping/Relay, do the following:

STEP 1 Click Configuration > IP Network Operations > DHCP > DHCP Snooping.

STEP 2 To enable DHCP Relay or DHCP Snooping enter the following fields:

- DHCP Snooping—Select to enable DHCP Snooping.
- Option 82 Passthrough—Select to leave foreign Option 82 information when forwarding packets.
- Verify MAC Address—Select to verify that the source MAC address of the Layer 2 header matches the client hardware address as appears in the DHCP Header (part of the payload) on DHCP untrusted ports.
- Backup Database—Select to back up the DHCP Snooping Binding database on the device's flash memory.

DHCP Interfaces

DHCP Relay and Snooping can be enabled on any interface with an IP Address, and on VLANs with or without an IP address.

To enable DHCP Snooping/Relay on specific interfaces, do the following:

STEP 1 Click Configuration > IP Network Operations > DHCP > DHCP Interfaces.

The following fields are displayed for each interface for which the features are enabled:

- Interface—On which DHCP Snooping/Relay is enabled or disabled.
- Interface IP Address—IP address of the interface on which DHCP Snooping/Relay is enabled.

- DHCP Snooping—Select to enable DHCP snooping.
- DHCP Relay— Select to enable DHCP Relay.

STEP 2 To enable DHCP Relay or DHCP Snooping on an interface, click ADD.

STEP 3 Select the interface and the features to be enabled: DHCP Relay or DHCP Snooping.

STEP 4 Click Apply. The settings are written to the Running Configuration file.

DHCP Snooping Trusted Interface

Packets from untrusted ports/LAGs are checked against the DHCP Snooping

Binding database (see the DHCP Snooping Binding Database page). By default, interfaces are untrusted.

To designate an interface as untrusted go to ARP Inspection.

IP Source Guard

IP Source Guard is a security feature that can be used to prevent traffic attacks caused when a host tries to use the IP address of its neighbor.

When IP Source Guard is enabled, the device only transmits client IP traffic to IP addresses contained in the DHCP Snooping Binding database. This includes both addresses added by DHCP Snooping and manually added entries.

If the packet matches an entry in the database, the device forwards it. If not, it is dropped.

Interactions with Other Features

The following points are relevant to IP Source Guard:

- DHCP Snooping must be globally enabled in order to enable IP Source Guard on an interface.
- IP source guard can be active on an interface only if the following apply:
- DHCP Snooping is enabled on at least one of the port's VLANs
- The interface is DHCP untrusted. All packets on trusted ports are forwarded.
- If a port is DHCP trusted, filtering of static IP addresses can be configured, even though IP Source Guard is not active in that condition by enabling IP Source Guard on the port.
- When the port's status changes from DHCP untrusted to DHCP trusted, the static IP address filtering entries remain in the Binding database, but they become inactive.
- Port security cannot be enabled if source IP and MAC address filtering is configured on a port.
- IP Source Guard uses TCAM resources and requires a single TCAM rule per IP Source Guard address entry. If the number of IP Source Guard entries exceeds the number of available TCAM rules, the extra addresses are inactive.

Filtering

If IP Source Guard is enabled on a port then the following apply:

- DHCP packets allowed by DHCP Snooping are permitted.

If source IP address filtering is enabled the following apply:

- IPv4 traffic—Only traffic with a source IP address that is associated with the port is permitted.
- Non IPv4 traffic—Permitted (Including ARP packets).

Configuring IP Source Guard Work Flow

To configure IP Source Guard:

STEP 1 Enable DHCP Snooping in the IP Network Operations > DHCP > DHCP Snooping page.

STEP 2 Define the VLANs on which DHCP Snooping is enabled in the IP Network Operations > DHCP > DHCP Interfaces page.

STEP 3 Configure interfaces as trusted or untrusted in the IP Network Operations > DHCP> Interface IP Settings page.

STEP 4 Enable IP Source Guard in the IP Network Operations > IP Source Guard > Feature Configuration page.

STEP 5 Enable IP Source Guard on the untrusted interfaces as required in the IP Network Operations > IP Source Guard > Interface Settings page.

STEP 6 View entries to the Binding database in the IP Network Operations > DHCP > DHCP Snooping Binding Database page.

IP Source Guard Interfaces

If IP Source Guard is enabled on an untrusted port orLAG, then DHCP packets allowed by DHCP Snooping are transmitted. If source IP address filtering is enabled, packet transmission is permitted as follows:

- IPv4 traffic — Only IPv4 traffic with a source IP address that is associated with the specific port is permitted.
- Non IPv4 traffic — All non-IPv4 traffic is permitted.

See Interactions with Other Features for more information about enabling IP Source Guard on interfaces.

To configure IP Source Guard on interfaces, use the Interface IP Settings page.

Binding Database

IP Source Guard uses the DHCP Snooping Binding database to check packets from untrusted ports. If the device attempts to write too many entries to the DHCP Snooping Binding database, the excessive entries are maintained in an inactive status. Entries are deleted when their lease time expires, and inactive entries may be made active.

See DHCP.

NOTE:
The Binding Database page only displays the entries in the DHCP Snooping Binding database defined on IP-Source-Guard-enabled ports.

To view the DHCP Snooping Binding database, use the DHCP Snooping Binding Database page.

To designate an interface as untrusted go to Interface IP Settings page.

DHCP Snooping Binding Database

See How the DHCP Snooping Binding Database is Built for a description of how dynamic entries are added to the DHCP Snooping Binding database.

Note the following points about maintenance of the DHCP Snooping Binding database:

- The device does not update the DHCP Snooping Binding database when a station moves to another interface.
- If a port is down, the entries for that port are not deleted.
- When DHCP Snooping is disabled for a VLAN, the binding entries that were collected for that VLAN are removed.
- If the database is full, DHCP Snooping continues to forward packets, but new entries are not created.

To add entries to the DHCP Snooping Binding database, do the following:

STEP 1 Click Configuration > IP Network Operations > DHCP Snooping Binding Database.

To see a subset of entries in the DHCP Snooping Binding database, enter the relevant search criteria and click Search.

The fields in the DHCP Snooping Binding Database are displayed.

These are described in the Add page, except for the IP Source Guard field:

- Status—
  - Active—IP Source Guard is active on the device.
  - Inactive—IP Source Guard is not active on the device.
- Reason—
  - No Problem
  - No Resource
  - No Snoop VLAN
  - Trust Port

STEP 2 To add an entry, click Add and enter the fields:

- VLAN ID—VLAN on which a packet is expected.
- MAC Address—MAC address of a packet.
- IPv4 Address—IP address of a packet.

Bindings Settings

- Interface—Type of interface on which a packet is expected.
- Type—The possible field values are the following:
  - Dynamic—Entry has limited lease time.
  - Static—Entry was statically configured.
- Lease Time—If the entry is dynamic, enter the amount of time that the entry is to be active in the DHCP Database in User Defined. If there is no Lease Time, check Infinite.)

STEP 3 Click Apply. The settings are defined, and the device is updated.

ARP Inspection

ARP enables IP communication within a Layer 2 Broadcast domain by mapping IP addresses to MAC addresses.

A malicious user can attack hosts, switches, and routers connected to a Layer 2 network by poisoning the ARP caches of systems connected to the subnet, and by intercepting traffic intended for other hosts on the subnet. This can happen because ARP allows a gratuitous reply from a host even if an ARP request was not received. After the attack, all traffic from the device under attack flows through the attacker's computer and then to the router, switch, or host.

The following shows an example of ARP cache poisoning.

ARP Cache Poisoning

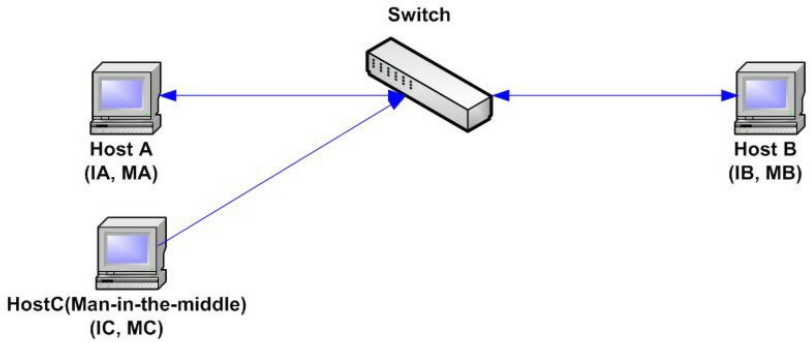

Hosts A, B, and C are connected to the switch on interfaces A, B and C, all of which are on the same subnet. Their IP and MAC addresses are shown in parentheses; for example, Host A uses IP address IA and MAC address MA. When Host A needs to communicate with Host B at the IP layer, it broadcasts an ARP request for the MAC address associated with IP address IB. Host B responds with an ARP reply. The switch and Host A update their ARP cache with the MAC and IP of Host B.

Host C can poison the ARP caches of the switch, Host A, and Host B by broadcasting forged ARP responses with bindings for a host with an IP address of IA (or IB) and a MAC address of MC. Hosts with poisoned ARP caches use the MAC address MC as the destination MAC address for traffic intended for IA or IB, which enables Host C to intercept that traffic. Because Host C knows the true MAC addresses associated with IA and IB, it can forward the intercepted traffic to those hosts by using the correct MAC address as the destination. Host C has inserted itself into the traffic stream from Host A to Host B, the classic man-in-the-middle attack.

How ARP Prevents Cache Poisoning

The ARP inspection feature relates to interfaces as either trusted or untrusted (see Security > ARP Inspection > Interface Setting page).

Interfaces are classified by the user as follows:

- Trusted — Packets are not inspected.
- Untrusted —Packets are inspected as described above.

ARP inspection is performed only on untrusted interfaces. ARP packets that are received on the trusted interface are simply forwarded.

Upon packet arrival on untrusted interfaces, the following logic is implemented:

- Search the ARP access control rules for the packet's IP/MAC addresses. If the IP address is found and the MAC address in the list matches the packet's MAC address, then the packet is valid; otherwise it is not.
- If the packet's IP address was not found, and DHCP Snooping is enabled for the packet's VLAN, search the DHCP Snooping Binding database for the packet's <VLAN - IP address> pair. If the <VLAN - IP address> pair was found, and the MAC address and the interface in the database match the packet's MAC address and ingress interface, the packet is valid.
- If the packet's IP address was not found in the ARP access control rules or in the DHCP Snooping Binding database the packet is invalid and is dropped. A SYSLOG message is generated.
- If a packet is valid, it is forwarded and the ARP cache is updated.

If the ARP Packet Validation option is selected (Properties page), the following additional validation checks are performed:

- Source MAC — Compares the packet's source MAC address in the Ethernet header against the sender's MAC address in the ARP request. This check is performed on both ARP requests and responses.
- Destination MAC — Compares the packet's destination MAC address in the Ethernet header against the destination interface's MAC address. This check is performed for ARP responses.
- IP Addresses — Compares the ARP body for invalid and unexpected IP addresses. Addresses include 0.0.0.0, 255.255.255.255, and all IP Multicast addresses.

Packets with invalid ARP Inspection bindings are logged and dropped. Up to 1024 entries can be defined in the ARP Access Control table.

Interaction Between ARP Inspection and DHCP Snooping

If DHCP Snooping is enabled, ARP Inspection uses the DHCP Snooping Binding database in addition to the ARP access control rules. If DHCP Snooping is not enabled, only the ARP access control rules are used.

ARP Defaults

The following table describes the ARP defaults:

Option	Default State
Dynamic ARP Inspection	Not enabled.
ARP Packet Validation	Not enabled.
ARP Inspection Enabled on VLAN	Not enabled.
Log Buffer Interval	SYSLOG message generation for dropped packets is enabled at 5 seconds interval

ARP Inspection Work Flow

To configure ARP Inspection, do the following:

STEP 1 Enable ARP Inspection and configure various options in the the Configuration > IP Network Operation > ARP Inspection > Feature Configuration page.

STEP 2 Configure interfaces as ARP trusted or untrusted in the Configuration > IP Network Operation > ARP Inspection > ARP Inspection Interface page.

STEP 3 Add rules in the Configuration > IP Network Operation > ARP Inspection > ARP Access Control and ARP Access Control Rules pages.

STEP 4 Define the VLANs on which ARP Inspection is enabled and the Access Control Rules for each VLAN in the Configuration > IP Network Operation > ARP Inspection > VLAN ARP Inspection page.

Defining ARP Inspection Properties

To configure ARP Inspection:

STEP 1 Click IP Network Operation > ARP Inspection > Feature Configuration.

Enter the following fields:

- ARP Inspection—Select to enable ARP Inspection.
- ARP Packet Validation—Select to enable the following validation checks:
- Log Buffer Interval—Select one of the following options:

- Never—Disabled SYSLOG dropped packet messages.

- Retry Frequency—Enable sending SYSLOG messages for dropped packets. Entered the frequency with which the messages are sent.

STEP 2 Click Apply. The settings are defined, and the Running Configuration file is updated.

Defining Dynamic ARP Inspection Interfaces Settings

Packets from untrusted ports/LAGs are checked against the ARP Access Rules table and the DHCP Snooping Binding database if DHCP Snooping is enabled (see the DHCP Snooping Binding Database page).

By default, ports/LAGs are ARP Inspection untrusted.

To change the ARP trusted status of a port/LAG, see Interface IP Settings.

ARP Access Control

To add entries to the ARP Inspection table, do the following:

STEP 1 Click Network Operation > ARP Inspection > ARP Access Control.

STEP 2 To add an entry, click Add.

STEP 3 Create a new ACL by entering the following:

Enter New Access Control

- ARP Access Control Name—Enter a user-created name.

Enter New Rule—

- IP Address—IP address of packet.
- MAC Address—MAC address of packet.

STEP 4 Click Apply. The settings are defined, and the Running Configuration file is updated.

Defining ARP Access Control Rules

To add more rules to a previously created ARP Access Control group, do the following:

STEP 1 Click Network Operation > ARP Inspection > ARP Access Control Rules.

The currently defined access rules are displayed.

STEP 2 To add more rules to a group, click Add.

STEP 3 Select a Access Control Group and enter the fields:

- IP Address—IP address of packet.
- MAC Address—MAC address of packet.

STEP 4 Click Apply. The settings are defined, and the Running Configuration file is updated.

VLAN ARP Inspection

To enable ARP Inspection on VLANs and associate Access Control Groups with a

VLAN, do the following:

STEP 1 Click P Network Operation > ARP Inspection > VLAN ARP Inspection.

STEP 2 To enable ARP Inspection on a VLAN, move the VLAN from the Available VLANs

list to the ARP Inspection Enabled list.

STEP 3 To associate an ARP Access Control group with a VLAN, click Add. Select the

VLAN number and select a previously defined ARP Access Control Name group.

STEP 4 Click Apply. The settings are defined, and the Running Configuration file is updated.

Interface Settings

To configure trusted interfaces, do the following:

STEP 1 Click Configuration > IP Network Operation > Interface Settings.

The following fields are displayed for each interface on which DHCP Snooping is enabled:

- Interface—Interface identifier.
- DHCP Snooping Trusted Interface—Whether the interface is DHCP Snooping trusted.
- IP Source Guard—Whether IP Source Guard is enabled on the interface.

ARP Inspection Trusted Interface—Whether the interface is ARP Inspection trusted.

Chapter 12 Security

Security

This section describes device security and access control. The system handles various types of security.

This chapter covers the following sections:

- Security Management
- RADIUS
- Network Access Control
- Port Security
- Storm Control

Security Management

The default username/password is admin/admin.

You can assign authentication methods to the various management access methods, such as, Telnet, HTTP, and HTTPS. The authentication can be performed locally or on a RADIUS server.

User Access & Accounts

The User Access & Accounts page enables entering additional users that are permitted to access to the device (read-only or read-write) or changing the passwords of existing users.

User authentication occurs in the order that the authentication methods are selected. If the first authentication method is not available, the next selected method is used. For example, if the selected authentication methods are RADIUS and Local, and all configured RADIUS servers are queried in priority order and do not reply, the user is authenticated locally.

If an authentication method fails or the user has insufficient privilege level, the user is denied access to the device. In other words, if authentication fails at an authentication method, the device stops the authentication attempt; it does not continue and does not attempt to use the next authentication method.

After adding a user (as described below), the default user is removed from the system.

NOTE:
It is not permitted to delete all users. If all users are selected, the Delete button is disabled.

To add a new user:

STEP 1 Click Configuration > Security > Security Management > User Access & Accounts.

This page displays the users defined in the system. Enter the following fields:

- HTTP Service—Select to enable on the device.
- HTTP Server Port—Enter the port on which HTTP is enabled.
- HTTPS Service—Select to enable on the device.
- HTTPS Server Port—Enter the port on which HTTPS is enabled.
- Telnet—Select to enable on the device.
- Authentication Methods—Select the options that are enabled on the device:
- Local—Only local user accounts (username/password) are checked.
- RADIUS—Authentication is performed through the RADIUS server.
- RADIUS, Local—Authentication is performed through the RADIUS server. If the server is not available, then local authentication is performed.
- None—No authentication is performed.

STEP 2 Click Add to add a new user or click Edit to modify a user.

STEP 3 Enter the parameters.

- User Name—Enter a new username between 0 and 20 characters. UTF-8 characters are not permitted.
- Password—Enter a password (UTF-8 characters are not permitted).
- Confirm Password—Enter the password again.

STEP 4 Click Apply. The user is added to the Running Configuration file of the device.

RADIUS

Remote Authorization Dial-In User Service (RADIUS) servers provide a centralized 802.1X or MAC-based network access control. The device is a RADIUS client that can use a RADIUS server to provide centralized security.

An organization can establish a RADIUS server to provide centralized 802.1X or MAC-based network access control for all of its devices. In this way, authentication and authorization can be handled on a single server for all devices in the organization.

The device can act as a RADIUS client that uses the RADIUS server for the following services:

- Authentication—Provides authentication of regular and 802.1X users logging onto the device with usernames and user-defined passwords.
- Authorization—Performed at login. After the authentication session is completed, an authorization session starts using the authenticated username. The RADIUS server then checks user privileges.

Defaults

No default RADIUS server is defined by default.

Radius Configuration

To use a RADIUS server, do the following:

STEP 1 Open an account for the device on the RADIUS server.

STEP 2 Configure that server along with the other parameters in the RADIUS pages.

NOTE:
If more than one RADIUS server has been configured, the device uses the configured priorities of the available RADIUS servers to select the RADIUS server to be used by the device.

To set the RADIUS server parameters, do the following:

STEP 1 Click Configuration > Security > RADIUS.

STEP 2 Enter the default RADIUS parameters if required. These are applied to all RADIUS servers that you configure. If a value is not entered for a specific server (in the Add RADIUS page) the device uses the values in these fields.

- Source IPv4 —Select the device IPv4 source interface to be used in messages for communication with the RADIUS server.
- Source IPv6 —Select the device IPv6 source interface to be used in messages for communication with the RADIUS server.
- Retries—Enter the number of transmitted requests that are sent to the RADIUS server before a failure is considered to have occurred.
- Timeout for Reply—Enter the number of seconds that the device waits for an answer from the RADIUS server before retrying the query, or switching to the next server.
- Dead Time—Enter the number of minutes that elapse before a non-responsive RADIUS server is bypassed for service requests. If the value is 0, the server is not bypassed.
- Key—Enter the default key string used for authenticating and encrypting between the device and the RADIUS server. This key must match the key configured on the RADIUS server. A key string is used to encrypt communications by using MD5.

This overrides the default key string if one has been defined.

STEP 3 Click Apply. The RADIUS default settings for the device are updated in the Running Configuration file.To add a RADIUS server, click Add.

STEP 4 Enter the values in the fields for each RADIUS server. To use the default values entered in the RADIUS page, select Use Default.

- Add Server—Select whether to specify the RADIUS server by IP Address or Name.
- IP Version—Select the version of the IP address of the RADIUS server.
- IPv6 Address Type—Select the IPv6 address type (if IPv6 is used). The options are the following:
- Link Local—The IPv6 address uniquely identifies hosts on a single network link. A link local address has a prefix of FE80, is not routable, and can be used for communication only on the local network. Only one link local address is supported. If a link local address exists on the interface, this entry replaces the address in the configuration.
- Global—The IPv6 address is a global Unicast IPV6 type that is visible and reachable from other networks.
- Interface—Select the link local interface (if IPv6 Address Type Link Local is selected) from the list.

- Server IP Address—Enter the IP address of the RADIUS server.
- Server Name—Enter the name of the RADIUS server.

Server Settings

- Authentication Port—Enter the UDP port number of the RADIUS server port for authentication requests.
- Priority—Enter the priority of the server. The priority determines the order the device attempts to contact the servers to authenticate a user. The device starts with the highest priority RADIUS server first. Zero is the highest priority. Key String—Enter the key string used for authenticating and encrypting communication between the device and the RADIUS server. This key must match the key configured on the RADIUS server.
- Usage Type—Enter the RADIUS server authentication type. The options are the following:
- Login—RADIUS server is used for authenticating users that ask to administer the device.
- 802.1x—RADIUS server is used for 802.1x authentication.
- All—RADIUS server is used for authenticating user that ask to administer the device and for 802.1X authentication.

STEP 5 Click Apply. The RADIUS server definition is added to the Running Configuration file of the device.

Network Access Control

This section describes 802.1x configuration. It covers the following topics:

- Overview
- Feature Configuration
- Port Authentication
- Authentication Hosts

Overview

802.1x authentication restricts unauthorized clients from connecting to a LAN through publicity-accessible ports. 802.1x authentication is a client-server model. In this model, network devices have the following specific roles.

- Client or supplicant
- Authenticator
- Authentication server

This is described in the figure below:

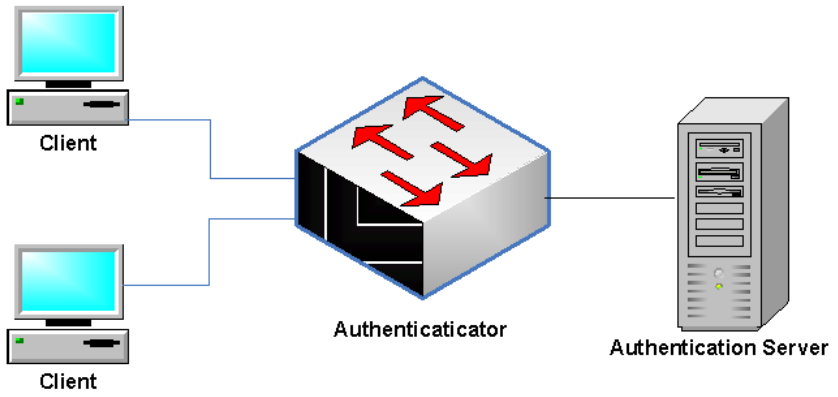

A network device can be either a client/supplicant, an authenticator or both per port.

Client or Supplicant

A client or supplicant is a network device that requests access to the LAN. The client is connected to an authenticator.

If the client uses the 802.1x protocol for authentication, it runs the supplicant part of the 802.1x protocol and the client part of the EAP protocol.

No special software is required on the client to use MAC-based authentication.

Authenticator

An authenticator is a network device that provides network services and to which supplicant ports are connected.

The following authentication modes on ports are supported:

- Multiple Host (802.1x)—Supports port-based authentication. If one client is authenticated, all client devices attaching to the port have access.
- Multiple Sessions—Supports client-based authentication. Each client must be authenticated individually before receiving access.

See Port Host Modes for more information.

The following authentication methods are supported:

- 802.1x-based—Supported in all authentication modes.
- MAC-based—Supported in all authentication modes.

In 802.1x-based authentication, the authenticator extracts the EAP messages from the 802.1x messages (EAPOL frames) and passes them to the authentication server, using the RADIUS protocol.

With MAC-based authentication, the authenticator itself executes the EAP client part of the software.

Authentication Server

An authentication server performs the actual authentication of the client. The authentication server for the device is a RADIUS authentication server with EAP extensions.

Port Administrative Authentication States

The port administrative state determines whether the client is granted access to the network.

The port administrative state can be configured in the Port Authentication page. The following values are available:

- Force Authorized Port authentication is disabled and the port transmits all traffic in accordance with its static configuration without requiring any authentication. The switch sends the 802.1x EAP-packet with the EAP success message inside when it receives the 802.1x EAPOL-start message. This is the default state.
- Force Unauthorized Port authentication is disabled and the port transmits all traffic via the guest VLAN and unauthenticated VLANs. For more information see Defining Host and Session Authentication. The switch sends 802.1x EAP packets with EAP failure messages inside when it receives 802.1x EAPOL- Start messages.
- Auto Enables 802.1 x authentications in accordance with the configured port host mode and authentication methods configured on the port.

Port Host Modes

Ports can be placed in the following port host modes (configured in the Host Authentication page):

- Multi-Host Mode

  A port is authorized if there is at least one authorized client.

  When a port is unauthorized and a guest VLAN is enabled, untagged traffic is remapped to the guest VLAN. Tagged traffic is dropped unless it belongs to the guest VLAN or to an unauthenticated VLAN. If guest VLAN is not enabled on a port, only tagged traffic belonging to unauthenticated VLANs is bridged.

  You can specify that untagged traffic from the authorized port will be remapped to a VLAN that is assigned by a RADIUS server during the authentication process. Tagged traffic is dropped unless it belongs to the RADIUS-assigned VLAN or to the unauthenticated VLANs. Radius VLAN assignment on a port is set in the Port Authentication page.

- Multi-Sessions Mode

  Unlike multi-host modes, a port in the multi-session mode does not have an authentication status. This status is assigned to each client connected to the port. This mode requires a TCAM lookup. Since Layer 3 mode switches (see Multi-Sessions Mode Support) do not have a TCAM lookup allocated for multi-sessions mode, they support a limited form of multi-sessions mode, which does not support guest VLAN and RADIUS VLAN attributes. The maximum number of authorized hosts allowed on the port is configured in the Port Authentication page.

Tagged traffic belonging to an unauthenticated VLAN is always bridged regardless of whether the host is authorized or not.

Tagged and untagged traffic from unauthorized hosts not belonging to an unauthenticated VLAN is remapped to the guest VLAN if it is defined and enabled on the VLAN, or it is dropped if the guest VLAN is not enabled on the port.

If an authorized host is assigned a VLAN by a RADIUS server, all its tagged and untagged traffic not belonging to the unauthenticated VLANs is bridged via the VLAN. If the VLAN is not assigned, all its traffic is bridged based on the static VLAN membership port configuration.

The LGS5xx in Layer 3 router mode supports the multi-sessions mode without guest VLAN and RADIUS-VLAN assignment:

Multiple Authentication Methods

If more than one authentication method is enabled on the switch, the following hierarchy of authentication methods is applied:

- 802.1x Authentication: Highest
- MAC-Based Authentication: Lowest

Multiple methods can run at the same time. When one method finishes successfully, the client becomes authorized, the methods with lower priority are stopped and the methods with higher priority continue.

When one of the authentication methods running simultaneously fails, the other methods continue.

When an authentication method finishes successfully for a client authenticated by a method with a lower priority, the attributes of the new method are applied. When the new method fails, the client is left authorized with the old method.

802.1x-Based Authentication

The device supports the 802.1x authentication mechanism, as described in the standard, to authenticate and authorize 802.1x supplicants.

The 802.1x-based authenticator relays transparent EAP messages between

802.1x supplicants and authentication servers. The EAP messages between supplicants and the authenticator are encapsulated into the 802.1x messages, and the EAP messages between the authenticator and authentication servers are encapsulated into the RADIUS messages.

This is described in the following:

Figure 1 802.1x-Based Authentication

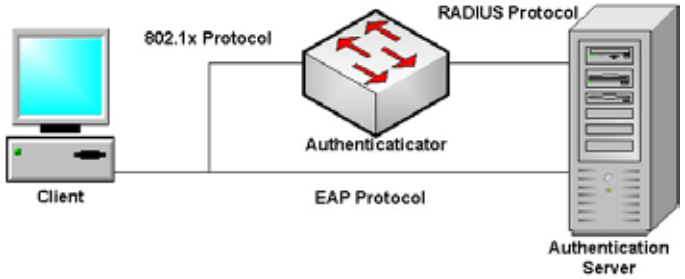

MAC-Based Authentication

MAC-based authentication is an alternative to 802.1X authentication that allows network access to devices (such as printers and IP phones) that do not have the 802.1X supplicant capability. MAC-based authentication uses the MAC address of the connecting device to grant or deny network access.

In this case, the switch supports EAP MD5 functionality with the username and password equal to the client MAC address, as shown below.

Figure 2 MAC-Based Authentication

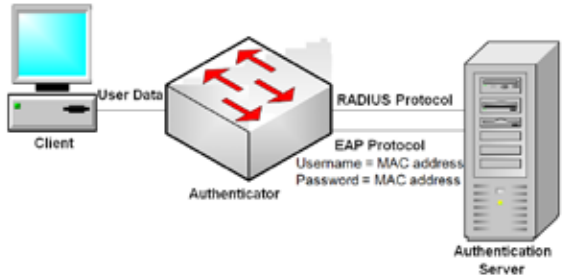

The method does not have any specific configuration.

Unauthenticated VLANs and the Guest VLAN

Unauthenticated VLANs and the guest VLAN provide access to services that do not require the subscribing devices or ports to be 802.1X or MAC-based authenticated and authorized.

The guest VLAN is the VLAN that is assigned to an unauthorized client. You can configure the guest VLAN and one or more VLANs to be unauthenticated in the Security > Network Access Control > Feature Configuration page.

An unauthenticated VLAN is a VLAN that allows access by authorized and unauthorized devices or ports.

An unauthenticated VLAN has the following characteristics:

- It must be a static VLAN, and cannot be the guest VLAN or the default VLAN.
- The member ports must be manually configured as tagged members.
- The member ports must be trunk ports. An access port cannot be member of an unauthenticated VLAN.

The guest VLAN, if configured, is a static VLAN with the following characteristics:

- It must be manually defined from an existing static VLAN.
- The guest VLAN cannot be used as the Voice VLAN or an unauthenticatedVLAN.

See "Table 3 Guest VLAN Support and RADIUS-VLAN Assignment Support" for a summary of the modes in which guest VLAN is supported.

Host Modes with Guest VLAN

The host modes work with guest VLAN in the following way:

- Single-Host and Multi-Host Mode

  Untagged traffic and tagged traffic belonging to the guest VLAN arriving on an unauthorized port are bridged via the guest VLAN. All other traffic is discarded. The traffic belonging to an unauthenticated VLAN is bridged via the VLAN.

- Multi-Sessions Mode in Layer 2

  Untagged traffic and tagged traffic, which does not belong to the unauthenticated VLANs and that arrives from unauthorized clients, is assigned to the guest VLAN using the TCAM rule and is bridged via the guest VLAN. The tagged traffic belonging to an unauthenticated VLAN is bridged via the VLAN.

  This mode cannot be configured on the same interface with policy-based VLANs.

- Multi-Sessions Mode in Layer 3

  The mode does not support the guest VLAN.

RADIUS VLAN Assignment or Dynamic VLAN AssignmentCommon Tasks

An authorized client can be assigned a VLAN by the RADIUS server, if this option is enabled in the Port Authentication page. This is called either Dynamic VLAN Assignment (DVA) or RADIUS-Assigned VLAN. In this guide, the term RADIUS Assigned VLAN is used.

When a port is in multi-session mode and RADIUS-Assigned VLAN is enabled, the device automatically adds the port as an untagged member of the VLAN that is assigned by the RADIUS server during the authentication process. The device classifies untagged packets to the assigned VLAN if the packets originated from the devices or ports that are authenticated and authorized.

NOTE:
In multi-session mode, RADIUS VLAN assignment is only supported when the device is in Layer 2 system mode.

When the RADIUS-Assigned VLAN feature is enabled, the host modes behave as follows:

- Single-Host and Multi-Host Mode

  Untagged traffic and tagged traffic belonging to the RADIUS-assigned VLAN are bridged via this VLAN. All other traffic not belonging to unauthenticated VLANs is discarded.

- Full Multi-Sessions Mode

  Untagged traffic and tagged traffic not belonging to the unauthenticated VLANs arriving from the client are assigned to the RADIUS-assigned VLAN using TCAM rules and are bridged via the VLAN.

- Multi-Sessions Mode in Layer 3 System Mode

  This mode does not support RADIUS-assigned VLAN.

The following table describes guest VLAN and RADIUS-assigned VLAN assignment support depending on authentication method and port mode.

Authenticati on Method	Single-host	Multi-host	Multi-sessions	
			Device in L3	Device in L2
802.1x	†	†	N/S	†
MAC	†	†	N/S	†

Legend:

†—The port mode supports the guest VLAN and RADIUS-VLAN assignment

N/S—The port mode does not support the authentication method.

Common Tasks

Workflow 1: To enable 802.1x authentication on a port:

STEP 1 Click Configuration > Security > Network Access Control > Feature Configuration.

STEP 2 Enable Port-based Authentication.

STEP 3 Select the Authentication Method.

STEP 4 Click Apply, and the Running Configuration file is updated.

STEP 5 Click Configuration > Security > Network Access Control > Port Authentication.

STEP 6 Select the required port and click Edit.

STEP 7 Set the Host Authentication mode.

STEP 8 Select a port, and click Edit.

STEP 9 Set the Administrative Port Control field to Auto.

STEP 10 Define the authentication methods.

STEP 11 Click Apply, and the Running Configuration file is updated.

Workflow 2: To configure 802.1x-based authentication

STEP 1 Click Configuration > Security > Network Control > > Port Authentication.

STEP 2 Select the required port and click Edit.

STEP 3 Enter the fields required for the port.The fields in this page are described in Port Authentication.

STEP 4 Click Apply, and the Running Configuration file is updated.

Use the Copy Settings button to copy settings from one port to another.

Workflow 4: To configure the guest VLAN:

STEP 1 Click Security > Network Access Control > Feature Configuration.

STEP 2 Select Enable in the Guest VLAN field.

STEP 3 Select the guest VLAN in the Guest VLAN ID field.

STEP 4 Click Apply, and the Running Configuration file is updated.

Workflow 5: To configure unauthenticated VLANs

STEP 1 Click Security > Network Access Control > Feature Configuration.

STEP 2 Select a VLAN, and click Edit.

STEP 3 Select a VLAN.

STEP 4 Optionally, uncheck Authentication to make the VLAN an unauthenticated VLAN.

STEP 5 Click Apply, and the Running Configuration file is updated.

Feature Configuration

The Feature Configuration page is used to globally enable 802.1X and define how ports are authenticated. For 802.1X to function, it must be activated globally and individually on each port.

To define port-based authentication:

STEP 1 Click Configuration > Security > Network Access Control > Feature Configuration.

STEP 2 Enter the parameters.

- Port-Based Authentication—Enable or disable port-based authentication.
- If this is disabled 802.1X is disabled.

If this is disabled 802.1X and MAC-based authentication is disabled.

- Authentication Method—Select the user authentication methods. The options are as follows:
- RADIUS, None—Perform port authentication first by using the RADIUS server. If no response is received from RADIUS (for example, if the server is down), then no authentication is performed, and the session is permitted. If the server is available but the user credentials are incorrect, access is denied and the session terminated.
- RADIUS—Authenticate the user on the RADIUS server. If no authentication is performed, the session is not permitted.

- None—Do not authenticate the user. Permit the session.
- Guest VLAN—Enable the use of a guest VLAN for unauthorized ports. If a guest VLAN is enabled, all unauthorized ports automatically join the VLAN selected in the Guest VLAN ID field. If a port is later authorized, it is removed from the guest VLAN.
- Guest VLAN ID—Select the guest VLAN from the list of VLANs.

The VLAN Authentication Table displays all VLANs, and indicates the authentication that has been enabled on them.

STEP 3 Click Edit to enable authentication on a VLAN.

STEP 4 Select the VLAN and select Authentication to enable authentication on the VLAN.

STEP 5 Click Apply. The settings are written to the Running Configuration file.

Port Authentication

The Port Authentication page enables configuration of 802.1X parameters for a port. Since some of the configuration changes are only possible while the port is in Force Authorized state, such as host authentication, it is recommended that you change the port control to Force Authorized before making changes. When the configuration is complete return the port control to its previous state.

NOTE:
A port with 802.1x defined on it cannot become a member of a LAG.

To configure 802.1X authentication:

STEP 1 Click Configuration > Security > Network Access Control > Port Authentication.

This page displays the authentication settings for all ports.

STEP 2 Select a port, and click Edit. Enter the following fields as follows:

- Port Control—Select the current port authorization state. If the state is Authorized, the port is either authenticated or the Administrative Port Control is Force Authorized. Conversely, if the state is Unauthorized, then the port is either not authenticated or the Administrative Port Control is Force Unauthorized. See Port Administrative Authentication States.
- Force Unauthorized— Port authentication is disabled and the port transmits all traffic via the guest VLAN and unauthenticated VLANs. For more information see Defining Host and Session Authentication.
- Auto—Enables 802.1 x authentications in accordance with the configured port host mode and authentication methods configured on the port.
- Force Authorized—Port authentication is disabled and the port transmits all traffic in accordance with its static configuration without requiring any authentication.
- Host Authentication Mode—Select one of the following options:
- Multiple Host (802.1x)—Supports port-based authentication with multiple clients per port.
- Multiple Sessions—Supports client-based authentication with multiple clients per port
- RADIUS VLAN Assignment—Select to enable Dynamic VLAN assignment on the selected port.
- Guest VLAN—Select to indicate that the usage of a previously-defined guest VLAN is enabled for the device.
- MAC 802.1x Based Authentication—Select to enable.

Port is authenticated based on the supplicant MAC address. Only 8 MAC-based authentications can be used on the port..

NOTE:
For MAC authentication to succeed, the RADIUS server supplicant username and password must be the supplicant MAC address. The MAC address must be in lower case letters and entered without the . or - separators; for example: 0020aa00bbcc.

- Periodic Reauthentication—Displays whether port reauthentication will be attempted after the specified Reauthentication Period.
- Reauthentication Period—Enter the number of seconds after which the selected port is reauthenticated.

STEP 3 Click Apply. The port settings are written to the Running Configuration file.

Authenticated Hosts

To display details about authenticated users, do the following:

STEP 1 Click Configuration > Security > Network Access Control > Authenticated Hosts.

This page displays the following fields:

- User Name—Supplicant names that were authenticated on each port.
- MAC Address—Displays the supplicant MAC address.
- Port—Number of the port.
- VLAN ID—Port's VLAN.
- Session Time—Amount of time that the supplicant was logged on the port.
- Authentication Method—Method by which the last session was authenticated.

Authentication Method and Port Mode Support

The following table shows which combinations of authentication method and port mode are supported.

Authentication Method	Multi-host	Multi-sessions	
		Device in L3	Device in L2
802.1x	†	†	†
MAC	†	†	†

Legend:

†—The port mode also supports the guest VLAN and RADIUS-VLAN assignment.
N/S—The authentication method does not support the port mode.

Mode Behavior

The following table describes how authenticated and non-authenticated traffic is handled in various situations.

	Unauthenticated Traffic				Authenticated Traffic			
	With Guest VLAN		Without Guest VLAN		With Radius VLAN		Without Radius VLAN	
	Untagged	Tagged	Untagged	Tagged	Untagged	Tagged	Untagged	Tagged
Single -host	Frames are re-mapped to the guest VLAN	Frames are dropped unless they belong to the guest VLAN or to the unauthenticated VLANs	Frames are dropped	Frames are dropped unless they belong to the unauthenticated VLANs	Frames are re-mapped to the RADIUS assigned VLAN	Frames are dropped unless they belong to the RADIUS VLAN or to the unauthenticated VLANs	Frames are bridged based on the static VLAN configuration	Frames are bridged based on the static VLAN configuration
Multi-host	Frames are re-mapped to the guest VLAN	Frames are dropped unless they belongs to the guest VLAN or to the unauthenticated VLANs	Frames are dropped	Frames are dropped unless they belongs to the unauthenticated VLANs	Frames are re-mapped to the Radius assigned VLAN	Frames are dropped unless they belongs to the Radius VLAN or to the unauthenticated VLANs	Frames are bridged based on the static VLAN configuration	Frames are bridged based on the static VLAN configuration
Lite multi-sessions	N/S	N/S	Frames are dropped	Frames are dropped unless they belongs to the unauthenticated VLANs	N/S	N/S	Frames are bridged based on the static VLAN configuration	Frames are bridged based on the static VLAN configuration
Full multi-sessions	Frames are re-mapped to the guest VLAN	Frames are re- mapped to the guest VLAN unless they belongs to the unauthenticated VLANs	Frames are dropped	Frames are dropped unless they belongs to the unauthenticated VLANs	Frames are re-mapped to the RADIUS assigned VLAN	Frames are re- mapped to the Radius VLAN unless they belongs to the unauthenticated VLANs	Frames are bridged based on the static VLAN configuration	Frames are bridged based on the static VLAN configuration

Port Security

Network security can be increased by limiting access on a port to users with specific MAC addresses. The MAC addresses can be either dynamically learned or statically configured.

Port security monitors received and learned packets. Access to locked ports is limited to users with specific MAC addresses.

Port Security has the following two modes:

- Classic Lock—All learned MAC addresses on the port are locked, and the port does not learn any new MAC addresses. The learned addresses are not subject to aging or relearning.
- Limited Dynamic Lock—The device learns MAC addresses up to the configured limit of allowed addresses. After the limit is reached, the device does not learn additional addresses. In this mode, the addresses are subject to aging and relearning.

When a frame from a new MAC address is detected on a port where it is not authorized (the port is classically locked, and there is a new MAC address, or the port is dynamically locked, and the maximum number of allowed addresses has been exceeded), the protection mechanism is invoked, and one of the following actions can take place:

- Frame is discarded
- Frame is forwarded
- Port is shut down

To configure port security do the following:

STEP 1 Click Configuration > Security > Port Security.

STEP 2 Select an interface to be modified, and click Edit.

STEP 3 Enter the parameters.

- Interface—Select the interface name.
- Interface Status—Select to lock the port.
- Learning Mode—Select the type of port locking. To configure this field, the Interface Status must be unlocked. The Learning Mode field is enabled only if the Interface Status field is locked. To change the Learning Mode, the Lock Interface must be cleared. After the mode is changed, the Lock Interface can be reinstated. The options are as follows:
- Classic Lock—Locks the port immediately, regardless of the number of addresses that have already been learned.
- Limited Dynamic Lock—Locks the port by deleting the current dynamic MAC addresses associated with the port. The port learns up to the maximum addresses allowed on the port. Both relearning and aging of MAC addresses are enabled.
- Maximum Addresses—Enter the maximum number of MAC addresses that can be learned on the port if Limited Dynamic Lock learning mode is selected. The number 0 indicates that only static addresses are supported on the interface.
- Action on Violation—Select an action to be applied to packets arriving on a locked port. The options are as follows:
- Discard—Discards packets from any unlearned source.
- Forward—Forwards packets from an unknown source without learning the MAC address.
- Shutdown—Discards packets from any unlearned source, and shuts down the port. The port remains shut down until reactivated, or until the device is rebooted.

STEP 4 Click Apply. Port security is modified, and the Running Configuration file is updated.

Storm Control

When Broadcast, Multicast, or Unknown Unicast frames are received, they are duplicated, and a copy is sent to all possible egress ports. This means that in practice they are sent to all ports belonging to the relevant VLAN. In this way, one ingress frame is turned into many, creating the potential for a traffic storm.

Storm protection enables you to limit the number of frames entering the device and to define the types of frames that are counted towards this limit.

When the rate of Broadcast, Multicast, or Unknown Unicast frames is higher than the user-defined threshold, frames received beyond the threshold are discarded.

To define Storm Control do the following:

STEP 1 Click Configuration > Security > Storm Control.

STEP 2 Select a port and click Edit.

STEP 3 Enter the parameters.

- Interface—Select the port for which storm control is enabled.
- Storm Control—Select to enable Storm Control.
- Storm Control Mode—Select one of the modes:
- Unknown Unicast, Multicast & Broadcast—Counts unknown Unicast, Broadcast, and Multicast traffic towards the bandwidth threshold.
- Multicast & Broadcast—Counts Broadcast and Multicast traffic towards the bandwidth threshold.
- Broadcast Only—Counts only Broadcast traffic towards the bandwidth threshold.
- Storm Control Rate Threshold—Enter the maximum rate at which unknown packets can be forwarded. The default for this threshold is 10,000 for FE devices and 100,000 for GE devices.

STEP 4 Click Apply. Storm control is modified, and the Running Configuration file is updated.

Chapter 13 Access Control List

Access Control List

The Access Control List (ACL) feature is part of the security mechanism.

ACLs enable network managers to define patterns (filter and actions) for ingress traffic. Packets, entering the device on a port or LAG with an active ACL, are either admitted or denied entry.

ACL definitions can also be used to define traffic flows in Quality of Service (QoS). For more information see Advanced Quality of Service.

This section contains the following topics:

- Access Control Lists
- Defining MAC-based ACLs
- IPv4-based ACLs
- Defining ACL Binding

Access Control Lists

An Access Control List (ACL) is an ordered list of classification filters and actions. Each single classification rule, together with its action, is called an Access Control Element (ACE).

Each ACE is made up of filters that distinguish traffic groups and associated actions. A single ACL may contain one or more ACEs, which are matched against the contents of incoming frames. Either a DENY or PERMIT action is applied to frames whose contents match the filter.

The device supports a maximum of 256 ACLs, and a maximum of 256 ACEs.

When a packet matches an ACE filter, the ACE action is taken and that ACL

processing is stopped. If the packet does not match the ACE filter, the next ACE is processed. If all ACEs of an ACL have been processed without finding a match, and if another ACL exists, it is processed in a similar manner.

NOTE:
If no match is found to any ACE in all relevant ACLs, the packet is dropped (as a default action). Because of this default drop action you must explicitly add ACEs into the ACL to permit the desired traffic, including management traffic, such as Telnet, HTTP or SNMP that is directed to the device itself. For example, if you do not want to discard all the packets that do not match the conditions in an ACL, you must explicitly add a lowest priority ACE into the ACL that permits all the traffic.

If IGMP snooping is enabled on a port bound with an ACL, add ACE filters in the ACL to forward IGMP/MLD packets to the device; otherwise, IGMP snooping fails at the port.

The order of the ACEs within the ACL is significant, since they are applied in a first- fit manner. The ACEs are processed sequentially, starting with the first ACE.

ACLs can be used for security, for example by permitting or denying certain traffic flows, and also for traffic classification and prioritization in the QoS Advanced mode.

NOTE:
A port can be either secured with ACLs or configured with advanced QoS policy, but not both.

There can only be one ACL per port.

To associate more than one ACL with a port, a policy with one or more class maps must be used.

The following types of ACLs can be defined (depending on which part of the frame header is examined):

- MAC ACL—Examines Layer 2 fields only, as described in Defining MAC-based ACLs
- IP ACL—Examines the Layer 3 layer of IP frames, as described in IPv4-based ACLs

If a frame matches the filter in an ACL, it is defined as a flow with the name of that ACL. In advanced QoS, these frames can be referred to using this Flow name, and QoS can be applied to these frames (see QoS Advanced Mode).

Creating ACLs Workflow

To create ACLs and associate them with an interface, perform the following:

1. Create one or more of the following types of ACLs:

a. MAC-based ACL by using the MAC Based ACL page and the MAC Based ACE page

b. IP-based ACL by using the IPv4 Based ACL page and the IPv4 Based ACE page.

2. Associate the ACL with interfaces by using the ACL Binding page.

Modifying ACLs Workflow

An ACL can only be modified if it is not in use. The following describes the process of unbinding an ACL in order to modify it:

1. If the ACL does not belong to a QoS Advanced Mode class map, but it has been associated with an interface, unbind it from the interface using the ACL Binding page.
2. If the ACL is part of the class map and not bound to an interface, then it can be modified.
3. If the ACL is part of a class map contained in a policy bound to an interface, you must perform the chain of unbinding as follows:

- Unbind the policy containing the class map from the interface by using Policy Binding.
- Delete the class map containing the ACL from the policy using the
- Configuring a Policy (Edit).
- Delete the class map containing the ACL, by using Defining Class Mapping.

Only then can the ACL be modified, as described in this section.

Defining MAC-based ACLs

MAC-based ACLs are used to filter traffic based on Layer 2 fields. MAC-based

ACLs check all frames for a match.

MAC-based ACLs are defined in the MAC Based ACL page. The rules are defined in the MAC Based ACE page.

To define a MAC-based ACL:

STEP 1 Click Configuration > Access Control List > MAC Based ACL. This page contains a list of all currently-defined MAC-based ACLs.

STEP 2 Click Add.

STEP 3 Enter the name of the new ACL in the ACL Name field. ACL names are case-sensitive.

STEP 4 Click Apply. The MAC-based ACL is saved to the Running Configuration file.

Adding Rules to a MAC-based ACL

To add rules (ACEs) to an ACL:

STEP 1 Click Configuration > Access Control List > MAC Based ACE.

STEP 2 Select an ACL, and click Search. The ACEs in the ACL are listed.

STEP 3 Click Add.

STEP 4 Enter the parameters.

- ACL Name—Select the name of the ACL to which an ACE is being added. ACE Settings
- ACE Priority—Enter the priority of the ACE. ACEs with higher priority are processed first. One is the highest priority.
- Action on Matched Packets—Select the action taken upon a match. The options are:
- Permit—Forward packets that meet the ACE criteria.
- Deny—Drop packets that meet the ACE criteria.
- Shutdown—Drop packets that meet the ACE criteria, and disable the port from where the packets were received. Such ports can be reactivated from the Port Settings page.

- Destination MAC Address—Select Any if all destination addresses are acceptable or User Defined to enter a destination address or a range of destination addresses.
- Destination MAC Address Value—Enter the MAC address to which the destination MAC address is to be matched and its mask (if relevant).
- Destination MAC Wildcard Mask—Enter the mask to define a range of MAC addresses. Note that this mask is different than in other uses, such as subnet mask. Here, setting a bit as 1 indicates don't care and 0 indicates to mask that value.

NOTE:
Given a mask of 0000 0000 0000 0000 0000 0000 1111 1111 (which means that you match on the bits where there is 0 and don't match on the bits where there are 1's). You need to translate the 1's to a decimal integer and you write 0 for each four zeros. In this example since 1111 1111 = 255, the mask would be written: as 0.0.0.255.

- Source MAC Address—Select Any if all source address are acceptable or User Defined to enter a source address or range of source addresses.
- Source MAC Address Value—Enter the MAC address to which the source MAC address is to be matched and its mask (if relevant).
- Source MAC Wildcard Mask—Enter the mask to define a range of MAC addresses.
- VLAN ID—Enter the VLAN ID section of the VLAN tag to match.
- 802.1p—Select Match to use 802.1p.
- 802.1p Value—Enter the 802.1p value to be added to the VPT tag.
- 802.1p Mask—Enter the wildcard mask to be applied to the VPT tag.
- EtherType—Enter the frame EtherType to be matched.

STEP 5 Click Apply. The MAC-based ACE is saved to the Running Configuration file.

IPv4-based ACLs

IPv4-based ACLs are used to check IPv4 packets, while other types of frames, such as ARPs, are not checked.

The following fields can be matched:

- IP protocol (by name for well-known protocols or directly by value)
- Source/destination ports for TCP/UDP traffic
- Flag values for TCP frames
- ICMP and IGMP type and code
- Source/destination IP addresses (including wildcards)
- DSCP/IP-precedence value

NOTE:
ACLs are also used as the building elements of flow definitions for per-flow QoS handling (see QoS Advanced Mode).

The IPv4 Based ACL page enables adding ACLs to the system. The rules are defined in the IPv4 Based ACE page.

Defining an IPv4-based ACL

To define an IPv4-based ACL:

STEP 1 Click Configuration > Access Control List > IPv4 Based ACL. This page contains all currently defined IPv4-based ACLs.

STEP 2 Click Add.

STEP 3 Enter the name of the new ACL in the ACL Name field. The names are case-sensitive.

STEP 4 Click Apply. The IPv4-based ACL is saved to the Running Configuration file.

Adding Rules (ACEs) to an IPv4-Based ACL

To add rules (ACEs) to an IPv4-based ACL:

STEP 1 Click Configuration > Access Control > IPv4-Based ACE.

STEP 2 Select an ACL, and click Search. All currently-defined IP ACEs for the selected ACL are displayed.

STEP 3 Click Add.

STEP 4 Enter the parameters.

- ACL Name—Displays the name of the ACL.

ACE Settings

- ACE Priority—Enter the priority. ACEs with higher priority are processed first.
- Action on Match Packets—Select the action assigned to the packet matching the ACE. The options are as follows:
- Permit—Forward packets that meet the ACE criteria.
- Deny—Drop packets that meet the ACE criteria.
- Shutdown—Drop packet that meets the ACE criteria and disable the port to which the packet was addressed. Ports are reactivated from the Port Management page.
- Protocol—Select to create an ACE based on a specific protocol or protocol ID. Select Any IPv4 to accept all IP protocols. Otherwise select one of the following protocols from the drop-down list:
- ICMP—Internet Control Message Protocol
- IGMP—Internet Group Management Protocol
- IP in IP—IP in IP encapsulation
- TCP—Transmission Control Protocol
- UDP—User Datagram Protocol
- Protocol ID —Instead of selecting the name, enter the protocol ID.
- Source IP Address—Select Any if all source address are acceptable or User Defined to enter a source address or range of source addresses.
- Source IP Address Value—Enter the IP address to which the source MAC address is to be matched and its mask (if relevant).
- Source IP Wildcard Mask—Enter the mask to define a range of IP addresses. Setting a bit as 1 indicates don't care and 0 indicates to mask that value.

NOTE:
Given a mask of 0000 0000 0000 0000 0000 0000 1111 1111 (which means that you match on the bits where there is 0 and don't match on the bits where there are 1's). You need to translate the 1's to a decimal integer and you write 0 for each four zeros. In this example since 1111 1111 = 255, the mask would be written: as 0.0.0.255.

- Destination IP Address—Select Any if all destination address are acceptable or User Defined to enter a destination address or range of destination addresses.
- Destination IP Address Value—Enter the IP address to which the destination IP address is to be matched.
- Destination IP Wildcard Mask—Enter the mask to define a range of IP addresses.
- Source Port—Select one of the following:
- Any—Match to all source ports.
- Single Port—Enter a single TCP/UDP source port to which packets are matched. This field is active only if 800/6-TCP or 800/17-UDP is selected in the Select from List drop-down menu.
- Destination Port—Select one of the available values that are the same as the Source Port field described above.

NOTE:
You must specify the IP protocol for the ACE before you can enter the source and/or destination port.

- Type of Service—The service type of the IP packet.
- Any—Any service type
- DSCP to Match—Differentiated Serves Code Point (DSCP) to match
- IP Precedence to match—IP precedence is a model of TOS (type of service) that the network uses to help provide the appropriate QoS commitments. This model uses the 3 most significant bits of the service type byte in the IP header, as described in RFC 791 and RFC 1349.

STEP 5 Click Apply. The IPv4-based ACE is saved to the Running Configuration file.

Defining ACL Binding

When an ACL is bound to an interface (port, LAG or VLAN), its ACE rules are applied to packets arriving at that interface. Packets that do not match any of the ACEs in the ACL are matched to a default rule, whose action is to drop unmatched packets.

multiple interfaces can be bound to the same ACL by grouping them into a policy-map, and binding that policy-map to the interface.

After an ACL is bound to an interface, it cannot be edited, modified, or deleted until it is removed from all the ports to which it is bound or in use.

NOTE:
It is possible to bind an interface (port, LAG or VLAN) to a policy or to an ACL, but they cannot be bound to both a policy and an ACL.

To bind an ACL to a port or LAG:

STEP 1 Click Configuration > Access Control List > ACL Binding (Port).

STEP 2 Select an interface type Ports/LAGs (Port or LAG).

STEP 3 Click Search. For each type of interface selected, all interfaces of that type are displayed with a list of their current ACLs:

NOTE:
To unbind all ACLs from an interface, select the interface, and click Clear.

STEP 4 Select an interface, and click Edit.

STEP 5 Select one of the following:

- MAC Based ACL—Select a MAC-based ACL to be bound to the interface.
- IPv4 Based ACL—Select an IPv4-based ACL to be bound to the interface.
- Permit Any Unmatched Packets—Select to enable/disable this action.

STEP 6 Click Apply. The ACL binding is modified, and the Running Configuration file is updated.

NOTE:
If no ACL is selected, the ACL(s) that is previously bound to the interface are unbound.

Chapter 14 Quality of Service

Quality of Service

The Quality of Service feature is applied throughout the network to ensure that network traffic is prioritized according to required criteria and the desired traffic receives preferential treatment.

This section covers the following topics:

- Overview
- Feature Configuration
- Queue Scheduling
- Bandwidth Control
- Basic QoS
- QoS Advanced Mode

Overview

The QoS feature is used to optimize network performance.

QoS provides the following:

- Classification of incoming traffic to traffic classes, based on attributes, includes the following:
- Device Configuration
- Ingress interface
- Packet content
- Combination of these attributes

QoS includes the following:

- Traffic Classification—Classifies each incoming packet as belonging to a specific traffic flow, based on the packet contents and/or the port.
- Assignment to Hardware Queues—Assigns incoming packets to forwarding queues. Packets are sent to a particular queue for handling as a function of the traffic class to which they belong. See Queue Scheduling.
- Other Traffic Class-Handling Attribute—Applies QoS mechanisms to various classes, including bandwidth management.

QoS Operation

The type of header field to be trusted is entered in the Global Settings page. For every value of that field, an egress queue is assigned, indicating through which queue the frame is sent (depending on whether the trust mode is CoS/802.1p or DSCP, respectively).

QoS Modes

The QoS mode that is selected applies to all interfaces in the system.

All traffic of the same class receives the same treatment, which is the single QoS action of determining the egress queue on the egress port, based on the indicated QoS value in the incoming frame. This can be the VLAN Priority Tag (VPT) 802.1p value in Layer 2 and the Differentiated Service Code Point (DSCP) value for IPv4 or Traffic Class (TC) value for IPv6 in Layer 3. When operating in Basic Mode, the device trusts this external assigned QoS value. The external assigned QoS value of a packet determines its traffic class and QoS.

The header field to be trusted is entered in the Global Settings page. For every value of that field, an egress queue is assigned where the frame is sent in the CoS/802.1p to Queue page or the DSCP to Queue page (depending on whether the trust mode is CoS/802.1p or DSCP, respectively).

- Advanced Mode—Per-flow Quality of Service (QoS).

  In Advanced Mode, a per-flow QoS consists of a class map and/or a policer:

  - A class map defines the kind of traffic in a flow, and contains one or more

    ACLs. Packets that match the ACLs belong to the flow.

  - A policer applies the configured QoS to a flow. The QoS configuration of a flow may consist of egress queue, the DSCP or CoS/802.1p value, and actions on out-of-profile (excess) traffic.
- Disable Mode—In this mode all traffic is mapped to a single best effort queue, so that no type of traffic is prioritized over another.

Only a single mode can be active at a time. When the system is configured to work in QoS Advanced Mode, settings for QoS Basic Mode are not active and vice versa.

When the mode is changed, the following occurs:

- When changing from QoS Advanced Mode to any other mode, policy profile definitions and class maps are deleted. ACLs bonded directly to interfaces remain bonded.
- When changing from QoS Basic Mode to Advanced Mode, the QoS Trust Mode configuration in Basic Mode is not retained.
- When disabling QoS, the shaper and queue setting (WRR/SP bandwidth setting) are reset to default values.

All other user configurations remain intact.

QoS Modes

The QoS mode that is selected applies to all interfaces in the system.

- Basic Mode—Class of Service (CoS).

  All traffic of the same class receives the same treatment, which is the single QoS action of determining the egress queue on the egress port, based on the indicated QoS value in the incoming frame. This can be the VLAN Priority Tag (VPT) 802.1p value in Layer 2 and the Differentiated Service Code Point (DSCP) value for IPv4 or Traffic Class (TC) value for IPv6 in Layer 3. When operating in Basic Mode, the device trusts this external assigned QoS value. The external assigned QoS value of a packet determines its traffic class and QoS.

  The header field to be trusted is entered in the Global Settings page. For every value of that field, an egress queue is assigned where the frame is sent in the CoS/802.1p to Queue page or the DSCP to Queue page (depending on whether the trust mode is CoS/802.1p or DSCP, respectively).
- Advanced Mode—Per-flow Quality of Service (QoS). In Advanced Mode, a per-flow QoS consists of a class map and/or a policer:
  - A class map defines the kind of traffic in a flow, and contains one or more ACLs. Packets that match the ACLs belong to the flow.
  - A policer applies the configured QoS to a flow. The QoS configuration of a flow may consist of egress queue, the DSCP or CoS/802.1p value, and actions on out-of-profile (excess) traffic.
- Disable Mode—In this mode all traffic is mapped to a single best effort queue, so that no type of traffic is prioritized over another.

Only a single mode can be active at a time. When the system is configured to work in QoS Advanced Mode, settings for QoS Basic Mode are not active and vice versa.

When the mode is changed, the following occurs:

- When changing from QoS Advanced Mode to any other mode, policy profile definitions and class maps are deleted. ACLs bonded directly to interfaces remain bonded.
- When changing from QoS Basic Mode to Advanced Mode, the QoS Trust Mode configuration in Basic Mode is not retained.

QoS Workflow

To configure general QoS parameters, perform the following:

STEP 1 Choose the QoS Mode (Basic, Advanced, or Disabled, as described in the "QoS Modes" section) for the system by using the QoS Properties page. The following steps in the workflow, assume that you have chosen to enable QoS.

STEP 2 Assign each interface a default CoS priority by using the QoS Feature Configuration page.

STEP 3 Assign the schedule method (Strict Priority or WRR) and bandwidth allocation for WRR to the egress queues by using the Queue page.

STEP 4 Designate an egress queue to each IP DSCP/TC value with the DSCP to Queue page. If the device is in DSCP trusted mode, incoming packets are put into the egress queues based on their DSCP/TC value.

STEP 5 Designate an egress queue to each CoS/802.1p priority. If the device is in CoS/802.1 Trusted Mode, all incoming packets are put into the designated egress queues according to the CoS/802.1p priority in the packets. This is done by using the CoS/802.1p to Queue page.

STEP 6 If required for Layer 3 traffic only, assign a queue to each DSCP/TC value, by using the DSCP to Queue page.

STEP 7 Enter bandwidth and rate limits in the following pages:

a Set egress shaping per queue by using the Egress Shaping Per Queue page.

b Set ingress rate limit and egress shaping rate per port by using the Bandwidth page.

STEP 8 Configure the selected mode by performing one of the following:

a Configure Basic Mode, as described in Workflow to Configure Basic QoS Mode

b Configure Advanced Mode, as described in Workflow to Configure Advanced QoS Mode.

Feature Configuration

The Feature Configuration Page contains fields for setting the QoS mode for the system (Basic, Advanced, or Disabled, as described in the "QoS Modes" section). In addition, the default CoS priority for each interface can be defined.

To select the QoS mode and configure QoS on an interface:

STEP 1 Click Configuration > Quality of Service > Feature Configuration.

STEP 2 Set the QoS mode. The following options are available:

- Disable — QoS is disabled on the device.
- Basic QoS — QoS is enabled on the device in Basic mode.
- AdvancedQoS—QoS is enabled on the device in Advanced Mode.

STEP 3 Select Port/LAG and click Search to display/modify all ports/LAGs on the device and their CoS information.

The following fields are displayed for all ports/LAGs:

- Interface — Type of interface.
- Default CoS—Default VPT value for incoming packets that do not have a VLAN Tag. The default CoS is 0. The default is only relevant for untagged frames and only if the system is in Basic Mode and Trust CoS is selected in the Global Settings page.

STEP 4 Click Apply. The Running Configuration file is updated.

STEP 5 Click Edit and enter the parameters.

- Interface — Select the port or LAG.
- Default CoS — Select the default CoS (Class-of-Service) value to be assigned for incoming packets (that do not have a VLAN tag).

STEP 6 Click Apply. The interface default CoS value is saved to Running Configuration file.

Queue Scheduling

The device supports 4 queues for each interface. Queue number four is the highest priority queue. Queue number one is the lowest priority queue.

There are two ways of determining how traffic in queues is handled, Strict Priority and Weighted Round Robin (WRR).

- Strict Priority — Egress traffic from the highest-priority queue is transmitted first. Traffic from the lower queues is processed only after the highest queue has been transmitted, thus providing the highest level of priority of traffic to the highest numbered queue.
- Weighted Round Robin (WRR)—In WRR mode the number of packets sent from the queue is proportional to the weight of the queue (the higher the weight the more frames are sent). For example, if there are a maximum of four queues possible and all four queues are WRR and the default weights are used, queue 1 receives 1/15 of the bandwidth (assuming all queues are saturated and there is congestion), queue 2 receives 2/15, queue 3 receives 4/15 and queue 4 receives 8 /15 of the bandwidth. The type of WRR algorithm used in the device is not the standard Deficit WRR (DWRR), but rather Shaped Deficit WRR (SDWRR).
- The queuing modes can be selected in the Queue Scheduling page. When the queuing mode is by strict priority, the priority sets the order in which queues are serviced, starting with Queue 4 (the highest priority queue) and going to the next lower queue when each queue is completed. When the queuing mode is Weighted Round Robin, queues are serviced until their quota has been used up and then another queue is serviced.
- It is also possible to assign some of the lower queues to WRR, while keeping some of the higher queues in strict priority. In this case traffic for the strict priority queues is always sent before traffic from the WRR queues. Only after the strict priority queues have been emptied is traffic from the WRR queues forwarded. (The relative portion from each WRR queue depends on its weight).

To select the priority method and enter WRR data, do the following.

STEP 1 Click Configuration > Quality of Service > Queue Scheduling.

STEP 2 Enter the parameters.

STEP 1 Click Configuration > Quality of Service > Queue Scheduling.

STEP 2 Enter the parameters.

- Queue—Displays the queue number.
- Scheduling Method
- Strict Priority—Traffic scheduling for the selected queue and all higher queues is based strictly on the queue priority.
- Weighted Round Robin Weight—If WRR is selected, enter the WRR
- weight assigned to the queue.
- %WRR Bandwidth—Displays the amount of bandwidth assigned to the queue. These values represent the percent of the WRR weight.

STEP 3 Click Apply. The queues are configured, and the Running Configuration file is updated.

Bandwidth Control

The Bandwidth Control page enables users to define two values, Ingress Rate Limit and Egress Shaping Rate, which determine how much traffic the system can receive and send.

The ingress rate limit is the number of bits per second that can be received from the ingress interface. Excess bandwidth above this limit is discarded.

The following values are entered for egress shaping:

- Committed Information Rate (CIR) sets the average maximum amount of data allowed to be sent on the egress interface, measured in bits per second
- Committed Burst Size (CBS) is the burst of data that is allowed to be sent, even though it is above the CIR. This is defined in number of bytes of data.

To enter bandwidth limitation, do the following:

STEP 1 Click Configuration > Quality of Service >Bandwidth Control.

The Bandwidth Control page displays bandwidth information for each interface.

STEP 2 Select an interface, and click Edit.

STEP 3 Select the Port or LAG interface.

STEP 4 Complete the following fields for the selected interface:

- Ingress Rate Control—Select to enable the ingress rate limit, which is defined in the field below.
- Ingress Rate Limit—Enter the maximum amount of bandwidth allowed on the interface.

NOTE The two Ingress Rate Limit fields do not appear when the interface type is LAG.

- Ingress Committed Burst Size—Enter the maximum burst size of data for the ingress interface in bytes of data. This amount can be sent even if it temporarily increases the bandwidth beyond the allowed limit. This field is only available if the interface is a port.
- Egress Shaping Control—Select to enable egress shaping on the interface.
- Egress Committed Information Rate—Enter the maximum bandwidth for the egress interface.
- Egress Committed Burst Size—Enter the maximum burst size of data for the egress interface in bytes of data. This amount can be sent even if it temporarily increases the bandwidth beyond the allowed limit.

STEP 5 Click Apply. The bandwidth settings are written to the Running Configuration file.

Basic QoS

In QoS Basic Mode, a specific domain in the network can be defined as trusted. Within that domain, packets are marked with 802.1p priority and/or DSCP to signal the type of service they require. Nodes within the domain use these fields to assign the packet to a specific output queue. The initial packet classification and marking of these fields is done in the ingress of the trusted domain.

Workflow to Configure Basic QoS Mode

To configure Basic QoS Mode, perform the following:

1. Select Basic Mode for the system by using the Feature Configuration page.
2. Select the trust-behavior using the Basic QoS page. The device supports CoS/802.1p Trusted Mode and DSCP Trusted Mode. CoS/802.1p Trusted Mode uses the 802.1p priority in the VLAN tag. DSCP Trusted Mode uses the DSCP value in the IP header.

If there is any port that, as an exception, should not trust the incoming CoS mark, disable the QoS state on that port using the Feature Configuration page.

Enable or disable the global selected trusted mode at the ports by using the Basic QoS page. If a port is disabled without trusted mode, all its ingress packets are forward in best effort with no guaranteed delivery. It is recommended that you disable the trusted mode at the ports where the CoS/802.1p and/or DSCP values in the incoming packets are not trustworthy. Otherwise, it might negatively affect the performance of your network

Configuration

The Basic QoS page contains information for enabling Trust on the device (see the Trust Mode field below). This configuration is active when the QoS mode is Basic Mode. Packets entering a QoS domain are classified at the edge of the QoS domain.

This page also enables configuring QoS on each port of the device, as follows:

QoS State Disabled on an Interface—All inbound traffic on the port is mapped to the best effort queue and no classification/prioritization takes place.QoS State of the Port is Enabled—Port prioritizes traffic on ingress based on the system wide configured trusted mode, which is either CoS/802.1p Trusted Mode or DSCP Trusted Mode.

To define the Trust configuration and configure QoS on interfaces:

STEP 1 Click Configuration > Quality of Service > QoS Basic Mode.

STEP 2 Select the Trust Mode while the device is in Basic Mode. If a packet CoS level and DSCP tag are mapped to separate queues, the trust mode determines the queue to which the packet is assigned:

- CoS/802.1p—Traffic is mapped to queues based on the VPT field in the VLAN tag, or based on the per-port default CoS/802.1p value (if there is no VLAN tag on the incoming packet).
- DSCP—All IP traffic is mapped to queues based on the DSCP field in the IPheader. If traffic is not IP traffic, it is mapped to the best effort queue.
- CoS/802.1p-DSCP—Either CoS/802.1p or DSCP whichever has been set.

STEP 3 Click Edit.

STEP 4 Select the Port or LAG interface.

STEP 5 Click to enable or disable QoS State for this interface.

STEP 6 Click Apply.

QoS Advanced Mode

Frames that match an ACL and were permitted entrance are implicitly labeled with the name of the ACL that permitted their entrance. Advanced Mode QoS actions can then be applied to these flows.

In QoS Advanced Mode, the device uses policies to support per-flow QoS. A

policy and its components have the following characteristics and relationships:

- A policy contains one or more class maps.
- A class map defines a flow with one or more associating ACLs. Packets that match only ACL rules (ACE) in a class map with Permit (forward) action are considered belonging to the same flow, and are subjected to the same quality of services. Thus, a policy contains one or more flows, each with a user-defined QoS.
- The QoS of a class map (flow) is enforced by the associating policer. There are two types of policers: single policer and aggregate policer. Each policer is configured with a QoS specification. A single policer applies the QoS to a single class map, and thus to a single flow, based on the policer QoS specification. An aggregate policer applies the QoS to one or more class maps, and thus one or more flows. An aggregate policer can support class maps from different policies.
- Per-flow QoS is applied to flows by binding the policies to the desired ports. A policy and its class maps can be bound to one or more ports, but each port is bound with at most one policy.

NOTE:

- Single policer and aggregation policer are available when the device is in Layer 2 mode.
- An ACL can be configured to one or more class maps regardless of policies.
- A class map can belong to only one policy.
- When a class map using single policer is bound to multiple ports, each port has its own instance of single policer; each applying the QoS on the class map (flow) at a port independent of each other.
- An aggregate policer applies the QoS to all its flow(s) in aggregation regardless of policies and ports.

Advanced QoS settings consist of three parts:

- Definitions of the rules to match. All frames matching a single group of rules are considered to be a flow.
- Definition of the actions to be applied to frames in each flow that match the rules.
- Binding the combinations of rules and action to one or more interfaces.

Workflow to Configure Advanced QoS Mode

1. Select Advanced Mode for the system by using the QoS Properties page.2.
2. Select the trust mode using the Global Settings page. If a packet CoS level and DSCP tag are mapped to separate queues, the trust mode determines the queue to which the packet is assigned:
- If internal DSCP values are different from those used on incoming packets, map the external values to internal values by using the Out-of-Profile DSCP Mapping page. This, in turn, opens the DSCP Remarking page.
3. Create ACLs as described in Create ACL Workflow.
4. If ACLs were defined, create class maps and associate the ACLs with them by using the Class Mapping page.
5. Create a policy using the Policy Table page, and associate the policy with one or more class maps using the Policy Class Map page. You can also specify the QoS, if needed, by assigning a policer to a class map when you associate the class map to the policy.
- Single Policer—Create a policy that associates a class map with a single policer by using the Policy Table page and the Class Mapping page. Within the policy, define the single policer.
- Aggregate Policer—Create a QoS action for each flow that sends all matching frames to the same policer (aggregate policer) by using the Aggregate Policer page. Create a policy that associates a class map with the aggregate policer by using the Policy Table page.
6. Bind the policy to an interface by using the Policy Binding page.

Configuring Global Settings

The Global Settings page contains information for enabling trust on the device. Packets entering a QoS domain are classified at the edge of the QoS domain.

To define the trust configuration:

STEP 1 Click Quality of Service > Advanced QoS.

STEP 2 Select the trusted mode while the device is in Advanced Mode. If a packet CoS level and DSCP tag are mapped to separate queues, the trusted mode determines the queue to which the packet is assigned:

- CoS/802.1p—Traffic is mapped to queues based on the VPT field in the VLAN tag, or based on the per-port default CoS/802.1p value (if there is no VLAN tag on the incoming packet), the actual mapping of the VPT to queue can be configured in the mapping CoS/802.1p to Queue page.
- DSCP—All IP traffic is mapped to queues based on the DSCP field in the IP header. The actual mapping of the DSCP to queue can be configured in the DSCP to Queue page. If traffic is not IP traffic, it is mapped to the best effort queue.
- CoS/802.1p-DSCP—Select to use Trust CoS Mode for non-IP traffic and Trust DSCP Mode for IP traffic.

STEP 3 Select the default Advanced Mode QoS trust mode (either Trusted or Untrusted) for interfaces in the Default Trust Mode field. This provides basic QoS functionality on Advanced QoS, so that you can trust CoS/DSCP on Advanced QoS by default (without having to create a policy).

In QoS Advanced Mode, when the Default Mode Status is set to Not Trusted, the Default CoS values configured on the interface are used for prioritizing the traffic arriving on the interface. See the Quality of Service > QoS Advanced Mode > Feature Configuration page for details.

If you have a policy on an interface then the Default Mode is irrelevant, the action is according to the policy configuration and unmatched traffic is dropped.

Defining Class Mapping

A Class Map defines a traffic flow with ACLs (Access Control Lists). A MAC ACL and an IP ACL can be combined into a class map. Class maps are configured to match packet criteria on a match-all or match-any basis. They are matched to packets on a first-fit basis, meaning that the action associated with the first-matched class map is the action performed by the system. Packets that match the same class map are considered to belong to the same flow.

NOTE:
Defining class maps does not have any effect on QoS; it is an interim step, enabling the class maps to be used later.

If more complex sets of rules are needed, several class maps can be grouped into a super-group called a policy (see Configuring a Policy).

The Class Mapping page shows the list of defined class maps and the ACLs comprising each, and enables you to add/delete class maps.

To define a Class Map:

STEP 1 Click Quality of Service > Advanced QoS > Class Mapping. This page displays the already-defined class maps.

STEP 2 Click Add.

A new class map is added by selecting one or two ACLs and giving the class map a name. If a class map has two ACLs, you can specify that a frame must match both ACLs, or that it must match either one or both of the ACLs selected.

STEP 3 Enter the parameters.

- Class Map Name—Enter the name of a new class map.
- Match ACLs—The criteria that a packet must match in order to be considered to belong to the flow defined in the class map. The options are the following:
  - IP—A packet must match either of the IP-based ACLs in the class map.
  - MAC—A packet must match the MAC-based ACL in the class map.
  - IP and MAC—A packet must match the IP-based ACL and the MAC-based ACL in the class map.
  - IP or MAC—A packet must match either the IP-based ACL or the MAC-based ACL in the class map.
- IP ACL—Select the IPv4-based ACL or the IPv6-based ACL for the class map.
- MAC ACL—Select the MAC-based ACL for the class map.
- Preferred ACL—Select whether packets are first matched to an IP-based ACL or a MAC-based ACL.

STEP 4 Click Apply. The Running Configuration file is updated.

Aggregate Policer

You can measure the rate of traffic that matches a predefined set of rules, and to enforce limits, such as limiting the rate of file transfer traffic that is allowed on a port.

This can be done by using the ACLs in the class map(s) to match the desired traffic, and by using a policer to apply the QoS on the matching traffic.

A policer is configured with a QoS specification. The following are the two kinds of policers:

- Single (Regular) Policer—A single policer applies the QoS to a single class map, and to a single flow based on the policer's QoS specification. When a class map using single policer is bound to multiple ports, each port has its own instance of single policer; each applying the QoS on the class map (flow) at ports that are otherwise independent of each other. A single policer is created in the Policy Table page.
- Aggregate Policer—An aggregate policer applies the QoS to one or more class maps, and one or more flows. An aggregation policer can support class maps from different policies. An aggregate policer applies QoS to all its flow(s) in aggregation regardless of policies and ports. An aggregate policer is created in the Aggregate Policer page.

  An aggregate policer is defined if the policer is to be shared with more than one class. Policers on a port cannot be shared with other policers in another device.

Each policer is defined with its own QoS specification with a combination of the following parameters:

- A maximum allowed rate, called a Committed Information Rate (CIR), measured in Kbps.
- An amount of traffic, measured in bytes, called a Committed Burst Size (CBS). This is traffic that is allowed to pass as a temporary burst even if it is above the defined maximum rate.

- An action to be applied to frames that are over the limits (called out-of-profile traffic), where such frames can be passed as is, dropped, or passed, but remapped to a new DSCP value that marks them as lower-priority frames for all subsequent handling within the device.

Assigning a policer to a class map is done when a class map is added to a policy. If the policer is an aggregate policer, you must create it using the Aggregate Policer page.

Defining Aggregate Policer

An aggregate policer applies the QoS to one or more class maps, therefore one or more flows. An aggregation policer can support class maps from different policies and applies the QoS to all its flow(s) in aggregation regardless of policies and ports.

NOTE:
The device supports aggregate policers and single policers only when operating in Layer 2 mode in devices that support a separate Layer 2 system mode.

STEP 1 Click Configuration > Quality of Service > Advanced QoS >Aggregate Policer. This page displays the existing aggregate policers.

STEP 2 Click Add.

STEP 3 Enter the parameters.

- Aggregate Policer Name—Enter the name of the Aggregate Policer.
- Ingress Committed Information Rate—Enter the maximum bandwidth allowed in bits per second. See the description of this in the Bandwidth page.
- Ingress Committed Burst Size—Enter the maximum burst size (even if it goes beyond the CIR) in bytes. See the description of this in the Bandwidth page.
- Exceed Action—Select the action to be performed on incoming packets that exceed the CIR. The following are possible values:
  - Forward—Packets exceeding the defined CIR value are forwarded.
  - Drop—Packets exceeding the defined CIR value are dropped.

STEP 4 Click Apply. The Running Configuration file is updated.

Policy Table

The Policy Table page displays the list of advanced QoS polices defined in the system. The page also allows you to create and delete polices. Only those policies that are bound to an interface are active (see Policy Binding page).

Each policy consists of the following:

- One or more class maps of ACLs which define the traffic flows in the policy.
- One or more aggregates that applies the QoS to the traffic flows in the policy.

After a policy has been added, class maps can be added by using the Policy Table page.

STEP 1 Click Configuration > Quality of Service> Policy Table. This page displays the list of defined policies.

STEP 2 Click Policy Class Map Table to display the Policy Class Maps page.

-or Click Add to open the Add Policy Table page.

STEP 3 Enter the name of the new policy in the Policy Name field.

STEP 4 Click Apply. The QoS policy profile is added, and the Running Configuration file is updated.

Policy Class Maps

One or more class maps can be added to a policy. A class map defines the type of packets that are considered to belong to the same traffic flow.

NOTE:
You cannot configure a policer to a class map when the device is operating in Layer 3 mode. The device supports policers only in Layer 2 mode.

To add a class map to a policy:

STEP 1 Click Configuration > Quality of Service > Policy Class Maps.

STEP 2 Select a policy in the Filter, and click Search. All class maps in that policy are displayed.

STEP 3 To add a new class map in Layer 2, click Add.

STEP 4 Enter the parameters.

- Policy Name—Displays the policy to which the class map is being added.
- Class Map Name—Select an existing class map to be associated with the policy. Class maps are created in the Class Mapping page.

Class Map Settings

- Policy Trust Mode—Select the action regarding the ingress CoS/802.1p and/or DSCP value of all the matching packets.
  - Use default trust mode—Ignore the ingress CoS/802.1p and/or DSCP value. The matching packets are sent as best effort.
  - Always Trust—If this option is selected, the device trusts the CoS/802.1p and DSCP of the matching packet. If a packet is an IP packet, the device puts the packet in the egress queue based on its DSCP value and the DSCP to Queue Table. Otherwise, the egress queue of the packet is based on the packet's CoS/802.1p value and the CoS/802.1p to Queue Table.
  - Set—If this option is selected, use the value entered in the New Value box to determine the egress queue of the matching packets as follows:
  - If the new value (0..7) is a CoS/802.1p priority, use the priority value and the CoS/802.1p to Queue Table to determine the egress queue of all the matching packets.
  - If the new value (0..63) is a DSCP, use the new DSCP and the DSCP to Queue Table to determine the egress queue of the matching IP packets.
  - Otherwise, use the new value (1..8) as the egress queue number for all the matching packets.
- Policer Type—Available in Layer 2 system mode only. Select the policer type for the policy. The options are as follows:
  - None—No policy is used.
  - Single Policer—The policer for the policy is a single policer.
  - Aggregate Policer—The policer for the policy is an aggregate policer.

If Policer Type is Single, enter the following QoS parameters:

- Ingress CIR—Enter the CIR in Kbps. See a description of this in the Bandwidth page.
- Ingress CBS—Enter the CBS in bytes. See a description of this in the Bandwidth page.
- Exceed Action—Select the action assigned to incoming packets exceeding the CIR. The options are:
  - None—No extra action, just forwarding.
  - Drop—Packets exceeding the defined CIR value are dropped.

STEP 5 To add a new class map in Layer 3, click Add.

STEP 6 Enter the parameters.

- Policy Name—Displays the policy to which the class map is being added.
- Class Map Name—Select an existing class map to be associated with the policy. Class maps are created in the Class Mapping page.

Class Map Settings

- Policy Trust Mode—Select the action regarding the ingress CoS/802.1p and/or DSCP value of all the matching packets.
  - Use default trust mode—Ignore the ingress CoS/802.1p and/or DSCP value. The matching packets are sent as best effort.
  - Always Trust—If this option is selected, the device trusts the CoS/802.1p and DSCP of the matching packet. If a packet is an IP packet, the device puts the packet in the egress queue based on its DSCP value and the DSCP to Queue Table. Otherwise, the egress queue of the packet is based on the packet's CoS/802.1p value and the CoS/802.1p to Queue Table.
  - Set—If this option is selected, use the value entered in the New Value box to determine the egress queue of the matching packets as follows:

    If the new value (0..7) is a CoS/802.1p priority, use the priority value and the CoS/802.1p to Queue Table to determine the egress queue of all the matching packets.

    If the new value (0..63) is a DSCP, use the new DSCP and the DSCP to Queue Table to determine the egress queue of the matching IP packets.

    Otherwise, use the new value (1..8) as the egress queue number for all the matching packets.

STEP 7 Click Apply.

Policy Binding

The Policy Binding page shows which policy profile is bound and to which port. When a policy profile is bound to a specific port, it is active on that port. Only one policy profile can be configured on a single port, but a single policy can be bound to more than one port.

When a policy is bound to a port, it filters and applies QoS to ingress traffic that belongs to the flows defined in the policy. The policy does not apply to traffic egress to the same port.

To edit a policy, it must first be removed (unbound) from all those ports to which it is bound.

NOTE:
It is possible to either bind a port to a policy or to an ACL but both cannot be bound.

To define policy binding, do the following:

STEP 1 Click Configuration > Quality of Service >Advanced QoS> Policy Binding.

STEP 2 Select a Policy Name and Interface Type if required.

STEP 3 Click Search. The policy is selected.

STEP 4 Select the following for the policy/interface:

- Bind—Select to bind the policy to the interface.
- Permit Any Unmatched Packets—Select to forward packets on the interface if they do not match any policy.

NOTE:
Permit Any can be defined only if IP Source Guard is not activated on the interface.

STEP 5 Click Apply. The QoS policy binding is defined, and the Running Configuration file is updated.

Port Policy

The Port Policy page displays the policy bound to an interface.

To display port policies:

STEP 1 Click Configuration > Quality of Service > Port Policy.

STEP 2 Select an Interface Type and Policy Name if required.

STEP 3 Click Search. The policy is selected and the Permit Any Unmatched Packets is displayed for the interface/policy.

Chapter 15 Maintenance

Maintenance

This section describes how to view system information and configure various options on the device.

It covers the following topics:

- Device Models
- System Mode & Reboot
- File Management
- Diagnostics

Device Models

All models can be fully managed through the web-based switch configuration utility. GE is the naming convention used for Gigabit Ethernet (10/100/1000) ports.

In Layer 2 system mode, the device forwards packets as a VLAN- aware bridge. In Layer 3 system mode, the device performs both IPv4 routing and VLAN-aware bridging.

NOTE:
Each model can be set to Layer 3 system mode by using the System Mode and Reboot page.

When the device operates in Layer 3 system mode, the QoS policers are not operational. Other QoS Advanced mode features are operational.

System Mode & Reboot

Some configuration changes, such as enabling jumbo frame support, require the system to be rebooted before they take effect. However, rebooting the device deletes the Running Configuration, so it is critical that the Running Configuration is saved to the Startup Configuration before the device is rebooted. Clicking Apply does not save the configuration to the Startup Configuration. For more information on files and file types, see the System Files section.

You can back up the configuration by using Maintenance > File Management > Configuration File Copy. You can also upload the configuration from a remote device. See the Configuration & Log Backup/Download (HTTP/HTTPS & Download) section.

- Reloading the device causes loss of connectivity in the network. By using delayed reboot, you can schedule the reboot to a time that is more convenient for users, e.g., late night.

To reboot the device, do the following:

STEP 1 Click Maintenance > Reboot.

STEP 2 Select Restore to Factory Defaults to reboot the device. This process erases the Startup Configuration file, and the backup configuration file.

- Reboot to Factory Defaults—Reboots the device by using the factory default configuration. This process erases the Startup Configuration file and the backup configuration file.

  The mirror configuration file is not deleted when restoring to factory defaults.

- System Mode—Select the Layer 2 or Layer 3 system mode.

STEP 3 Click Apply. The Running Configuration file is updated.

File Management

This section describes how system files are managed. The following topics are covered:

- Overview
- Firmware & Boot Code
- Active Firmware Image
- Configuration & Log
- Configuration File Copy

Overview

System files are files that contain configuration information, firmware images or boot code.

Various actions can be performed with these files, such as: selecting the firmware file from which the device boots, copying various types of configuration files internally on the device, or copying files to or from an external device, such as an external server.

The possible methods of file transfer are as follows:

- Internal copy.
- HTTP/HTTPS that uses the facilities that the browser provides.
- TFTF client, requiring a TFTP server.

Configuration files on the device are defined by their type, and contain the settings and parameter values for the device.

When a configuration is referenced on the device, it is referenced by its configuration file type (such as Startup Configuration or Running Configuration), as opposed to a file name that can be modified by the user.

Content can be copied from one configuration file type to another, but the names of the file types cannot be changed by the user.

Other files on the device include firmware, boot code, and log files, and are referred to as operational files.

The configuration files are text files and can be edited in a text editor, such as Notepad after they are copied to an external device, such as a PC.

Files and File Types

The following types of configuration and operational files are found on the device:

- Running Configuration — Contains the parameters currently being used by the device to operate. This is the only file type that is modified when you change parameter values on the device. If the device is rebooted, the Running Configuration is lost. The Startup Configuration, stored in flash memory, overwrites the Running Configuration, stored in RAM. To preserve any changes you made to the device, you must save the Running Configuration to the Startup Configuration or another file type.
- Startup Configuration — The parameter values that were saved by copying another configuration (usually the Running Configuration) to the Startup Configuration. The Startup Configuration is retained in flash memory and is preserved when the device is rebooted. At this time, the Startup Configuration is copied to RAM and identified as the Running Configuration.
- Backup Configuration — A manual copy of a configuration file used for protection against system shutdown or for the maintenance of a specific operating state. You can copy the Startup Configuration, or Running Configuration to a Backup Configuration file. The Backup Configuration exists in flash memory and is preserved if the device is rebooted.
- Firmware — The program that controls the operations and functionality of the device; more commonly referred to as the image.
- Boot Code — Controls the basic system startup and launches the firmware image.
- Flash Log — SYSLOG messages stored in Flash memory.

File Actions

The following actions can be performed to manage firmware and configuration files:

- Upgrade the firmware or boot code as described in Overview section.
- View the firmware image currently in use or select the image to be used in the next reboot as described in the Active Firmware Image section.
- Save configuration files on the device to a location on another device as described in the Configuration & Log section.
- Copy one configuration file type to another configuration file type as described in the Configuration File Copy section.

Firmware & Boot Code

The Upgrade/Backup Firmware process can be used to do the following:

- Upgrade or backup the firmware image.
- Upgrade or backup the boot code.

The following methods for transferring files are supported:

- HTTP/HTTPS that uses the facilities provided by the browser
- TFTP that requires a TFTP server

When you upgrade the firmware and reboot the device the new firmware loads.To upgrade or backup a software image, do the following:

STEP 1 Click Maintenance > File Management > Firmware & Boot Code.

STEP 2 Select the Transfer Method. Proceed as follows:

- If you selected TFTP, go to **STEP 3**.
- If you selected HTTP/HTTPS, go to **STEP 4**.

STEP 3 If you selected via TFTP, enter the parameters as described in this step. Otherwise, skip to **STEP 4**.

Select one of the following options for Command:

- Upgrade—Specifies that the file type on the device is to be replaced with a new version of that file type located on a TFTP server.
- Backup—Specifies that a copy of the file type is to be saved to a file on another device.

Enter the following fields:

- File Type—Select the destination file type:
- Firmware—The program that controls the operations and functionality of the device; more commonly referred to as the image.
- Boot Code—Controls the basic system startup and launches the firmware image.
- Source File Name—Enter the name of the source file.
- TFTP Server—Select whether to specify the TFTP server by IP address or domain name.
- IP Version—Select whether an IPv4 or an IPv6 address is used.
- IPv6 Address Type—Select the IPv6 address type (if IPv6 is used). The options are as follows:
- Link Local—The IPv6 address uniquely identifies hosts on a single network link. A link local address has a prefix of FE80, is not routable, and can be used for communication only on the local network. Only one link local address is supported. If a link local address exists on the interface, this entry replaces the address in the configuration.
- Global—The IPv6 address is a global Unicast IPV6 type that is visible and reachable from other networks.
- Interface—Select the link local interface (if IPv6 is used) from the list.
- TFTP Server IP Address—Enter the IP address of the TFTP server.
- TFTP Server Name—Enter the domain name of the TFTP server.

STEP 4 If you selected via HTTP/HTTPS, you can only Upgrade. Enter the parameters as described in this step.

- File Type—Select Firmware Image to upgrade the firmware image.
- Source File Name—Click Browse to select a file or enter the path and source file name to be used in the transfer.

STEP 5 Click Apply.

NOTE:
When the process is completed, the following information is displayed:

- Bytes Transferred—How many bites were transferred in the process.
- Status—Did the process succeed or fail.
- Error Message—Reason for failure of the process.

Active Firmware Image

There are two firmware images stored on the device. One of the images is identified as the active image and the other image is identified as the inactive image. The device boots from the image you set as the active image. You can change the image identified as the inactive image to the active image.

To select the active image, do the following:

STEP 1 Click Maintenance > File Management > Active Firmware Image.

The page displays the following:

- Active Firmware Image—Displays the image file that is currently active on the device.
- Version —Displays the firmware version of the active image.
- Active Firmware Image After Reboot—Displays the image that is active after reboot.
- Version —Displays the firmware version of the active image as it will be after reboot.

STEP 2 Select the image from the Active Firmware Image After Reboot menu to identify the firmware image that is used as the active image after the device is rebooted. The version number associated with it displays the firmware version of the active image that is used after the device is rebooted.

STEP 3 Click Apply. The active image selection is updated.

Configuration & Log

The Configuration & Log (Backup & Download) page enables the following:

- Backing up configuration files or logs from the device to an external device.
- Restoring configuration files from an external device to the device.When restoring a configuration file to the Running Configuration, the imported file adds any configuration commands that did not exist in the old file and overwrites any parameter values in the existing configuration commands.

When restoring a configuration file to the Running Configuration, the imported file adds any configuration commands that did not exist in the old file and overwrites any parameter values in the existing configuration commands.

When restoring a configuration file to the Startup Configuration or a backup configuration file, the new file replaces the previous file.

When restoring to Startup Configuration, the device must be rebooted for the restored Startup Configuration to be used as the Running Configuration. You can reboot the device by using the process described in the Management Interface section.

To backup or restore the system configuration file, do the following:

STEP 1 Click Maintenance > File Management > Configuration & Log.

STEP 2 Select the File Transfer Protocol.

STEP 3 If you selected via TFTP, enter the parameters. Otherwise, skip to **STEP 5**.

Enter the following fields:

- Command—Select one of the following options:
- Download—Specifies that the file on another device upgrades a file type on the device.
- Backup—Specifies that a file type is to be copied to a file on another device.
- Source File Name—Enter the source file name for download. File names cannot contain slashes (\ or /), cannot start with a period (.), and must include between 1 and 160 characters. (Valid characters: A-Z, a-z, 0-9, “.”, “-”,“_”).
- Destination File —Select one of the files displayed as the file to be upgraded. Only valid file types are displayed. (The file types are described in the Files and File Types section).
- TFTP Server—Select whether to specify the TFTP server by IP address or domain name.
- IP Version—Select whether an IPv4 or an IPv6 address is used.
- IPv6 Address Type—Select the IPv6 address type (if IPv6 is used). The options are:
- Link Local—The IPv6 address uniquely identifies hosts on a single network link. A link local address has a prefix of FE80, is not routable, and can be used for communication only on the local network. Only one link local address is supported. If a link local address exists on the interface, this entry replaces the address in the configuration.
- Global—The IPv6 address is a global Unicast IPV6 type that is visible and reachable from other networks.
- Interface—Select the link local interface (if IPv6 is used) from the list.
- TFTP Server IP Address—Enter the IP address of the TFTP server.

- TFTP Server Name—Enter the domain name of the TFTP server.

NOTE:
If the server is selected by name in the Server Definition, there is no need to select the IP version-related options.

STEP 4 Click Apply. The file is upgraded or backed up.

STEP 5 If you selected via HTTP/HTTPS, enter the parameters as described in this step.

- Command—Select one of the following options:
- Download—Download a new version of a file (upgrade).
- Backup—Upload a file.
- Source File Name—Enter the file name for download.
- Destination File—Select the configuration file type to be downloaded to.
- Only valid file types are displayed. (The file types are described in the Files and File Types section).

STEP 6 Click Apply. The file is upgraded or backed up.

NOTE:
When the process initiated is completed, the following information is displayed:

- Bytes Transferred—How many bites were transferred in the process.
- Status—Did the process succeed or fail.
- Error Message—Reason for failure of the process.

Configuration File Copy

When you click Apply on any window, changes that you made to the device configuration settings are stored only in the Running Configuration. To preserve the parameters in the Running Configuration, the Running Configuration must be copied to another configuration type or saved on another device.

CAUTION Unless the Running Configuration is copied to the Startup Configuration or another configuration file, all changes made since the last time the file was copied are lost when the device is rebooted.

The following combinations of copying internal file types are allowed:

- From the Running Configuration to the Startup Configuration or Backup Configuration.
- From the Startup Configuration to the Running Configuration or Backup Configuration.
- From the Backup Configuration to the Running Configuration, Startup Configuration.

To copy one type of configuration file to another type of configuration file, do the following:

STEP 1 Click Maintenance > File Management > Configuration File Copy.

STEP 2 Select the Source File to be copied. Only valid file types are displayed (described in the Files and File Types section).

STEP 3 Select the Destination File to be overwritten by the source file.

STEP 4 Click Apply. The file is copied.

Diagnostics

This section covers the following topics:

- Optical Module Status
- Ping
- Port Mirroring

Optical Module Status

The Optical Module Status page displays the operating conditions reported by the SFP (Small Form-factor Pluggable) transceiver. Some information might not be available for SFPs that do not support the digital diagnostic monitoring standard SFF-8472.

To view the results of optical tests, click Maintenance > Diagnostics > Optical Module Status.

Module Status.

This page displays the following fields:

- Port—Port number on which the SFP is connected.
- Description—Description of optical transceiver.
- Serial Number—Serial number of optical transceiver.
- Data Ready—SFP is operational. Values are True and False
- Loss of Signal—Local SFP reports signal loss. Values are True and False.
- Transmitter Fault—Remote SFP reports signal loss. Values are True, False, and No Signal (N/S).
- Temperature—Temperature (Celsius) at which the SFP is operating.

Ping

Ping is a utility used to test if a remote host can be reached and to measure the round-trip time for packets sent from the device to a destination device.

Ping operates by sending Internet Control Message Protocol (ICMP) echo request packets to the target host and waiting for an ICMP response, sometimes called a pong. It measures the round-trip time and records any packet loss.

To ping a host, do the following:

STEP 1 Click Maintenance > Diagnostics > > Ping.

STEP 2 Configure ping by entering the fields:

- Target—Select whether to specify the source interface by its IP address or name. This field influences the interfaces that are displayed in the Source IP field, as described below.
- IP Version—If the source interface is identified by its IP address, select either IPv4 or IPv6 to indicate that it will be entered in the selected format.
- IPv6 Address Type—Select Link Local or Global as the type of IPv6 address to enter as the destination IP address.
- Link Local—The IPv6 address uniquely identifies hosts on a single network link. A link local address has a prefix of FE80, is not routable, and can be used for communication only on the local network. Only one link local address is supported. If a link local address exists on the interface, this entry replaces the address in the configuration.
- Global—The IPv6 address is a global Unicast IPV6 type that is visible and reachable from other networks.
- Interface—If the IPv6 address type is Link Local, select from where it is received.
- Target IP Address—Address of the device to be pinged. Whether this is an IP address or host name depends on the Host Definition.
- Target Name—Host name of the device to be pinged. Whether this is an IP address or host name depends on the Host Definition.
- Ping Interval—Length of time the system waits between ping packets. Ping is repeated the number of times configured in the Number of Pings field, whether the ping succeeds or not. Choose to use the default interval or specify your own value.
- Number of Pings—The number of times the ping operation is performed.

- Choose to use the default or specify your own value.

STEP 3 Click Start to ping the host. The ping status appears and another message is added to the list of messages, indicating the result of the ping operation.

STEP 4 View the results of ping in the Ping Result section of the page:

- Result—Success or fail of ping.
- Number of Pings Sent—Numbers of responses sent.
- Number of Ping Responses Received—Numbers of responses received.
- Packets Lost—Numbers of responses not received
- Minimum Round Trip Time—Minimum time passed between sending of packets and reception of responses.
- Maximum Round Trip Time—Maximum time passed between sending of packets and reception of responses
- Average Round Trip Time—Average time passed between sending of packets and reception of responses.

Port Mirroring

Port mirroring is used on a network device to send a copy of network packets seen on one or multiple device ports, to a network monitoring connection on another port on the device. This is commonly used for network appliances that require monitoring of network traffic, such as an intrusion-detection system. A network analyzer connected to the monitoring port processes the data packets for diagnosing, debugging, and performance monitoring. Up to four sources can be mirrored. This can be any combination of four individual ports.

A packet that is received on a network port assigned to a VLAN that is subject to mirroring is mirrored to the analyzer port even if the packet was eventually trapped or discarded. Packets sent by the device are mirrored when Transmit (Tx) mirroring is activated.

Mirroring does not guarantee that all traffic from the source port(s) is received on the analyzer (destination) port. If more data is sent to the analyzer port than it can support, some data might be lost.

Only one instance of mirroring is supported system-wide. The analyzer port is the same for all the mirrored ports.

To enable mirroring, do the following:

STEP 1 Click Maintenance > Diagnostics > Port Mirroring.

The following fields are displayed:

- Destination Port—Port to which traffic is to be copied; the analyzer port.
- Source Port—Interface, port, from which traffic is sent to the analyzer port.
- Mirror Type—Type of monitoring: incoming to the port (Rx), outgoing from the port (Tx), or both.
- Status— Displays one of the following values:
- Active—Both source and destination interfaces are up and forwarding traffic.
- Not Ready—Either source or destination (or both) are down or not forwarding traffic for some reason.

STEP 2 Click Add to add a port to be mirrored.

STEP 3 Enter the parameters:

- Destination Port—Select the analyzer port to where packets are copied. A network analyzer, such as a PC running Wireshark, is connected to this port. If a port is identified as an analyzer destination port, it remains the analyzer destination port until all entries are removed.
- Source Port—Select the source port from where traffic is to be mirrored.
- Mirror Type—Select whether incoming, outgoing, or both types of traffic are mirrored to the analyzer port. If Port is selected, the options are as follows:
- Rx Only—Port mirroring on incoming packets.
- Tx Only—Port mirroring on outgoing packets.
- Tx and Rx—Port mirroring on both incoming and outgoing packets.

STEP 4 Click Apply. Port mirroring is added to the Running Configuration.

Support

Click “Get Support” and take you to Linksys Small Business support website, where you can find out all the ways you can get support and help for your Linksys Small Business products, such as set-up guidance, top frequently asked questions, download software, live chat with our technical support, or community forum

Visit **linksys.com/support** for award-winning technical support

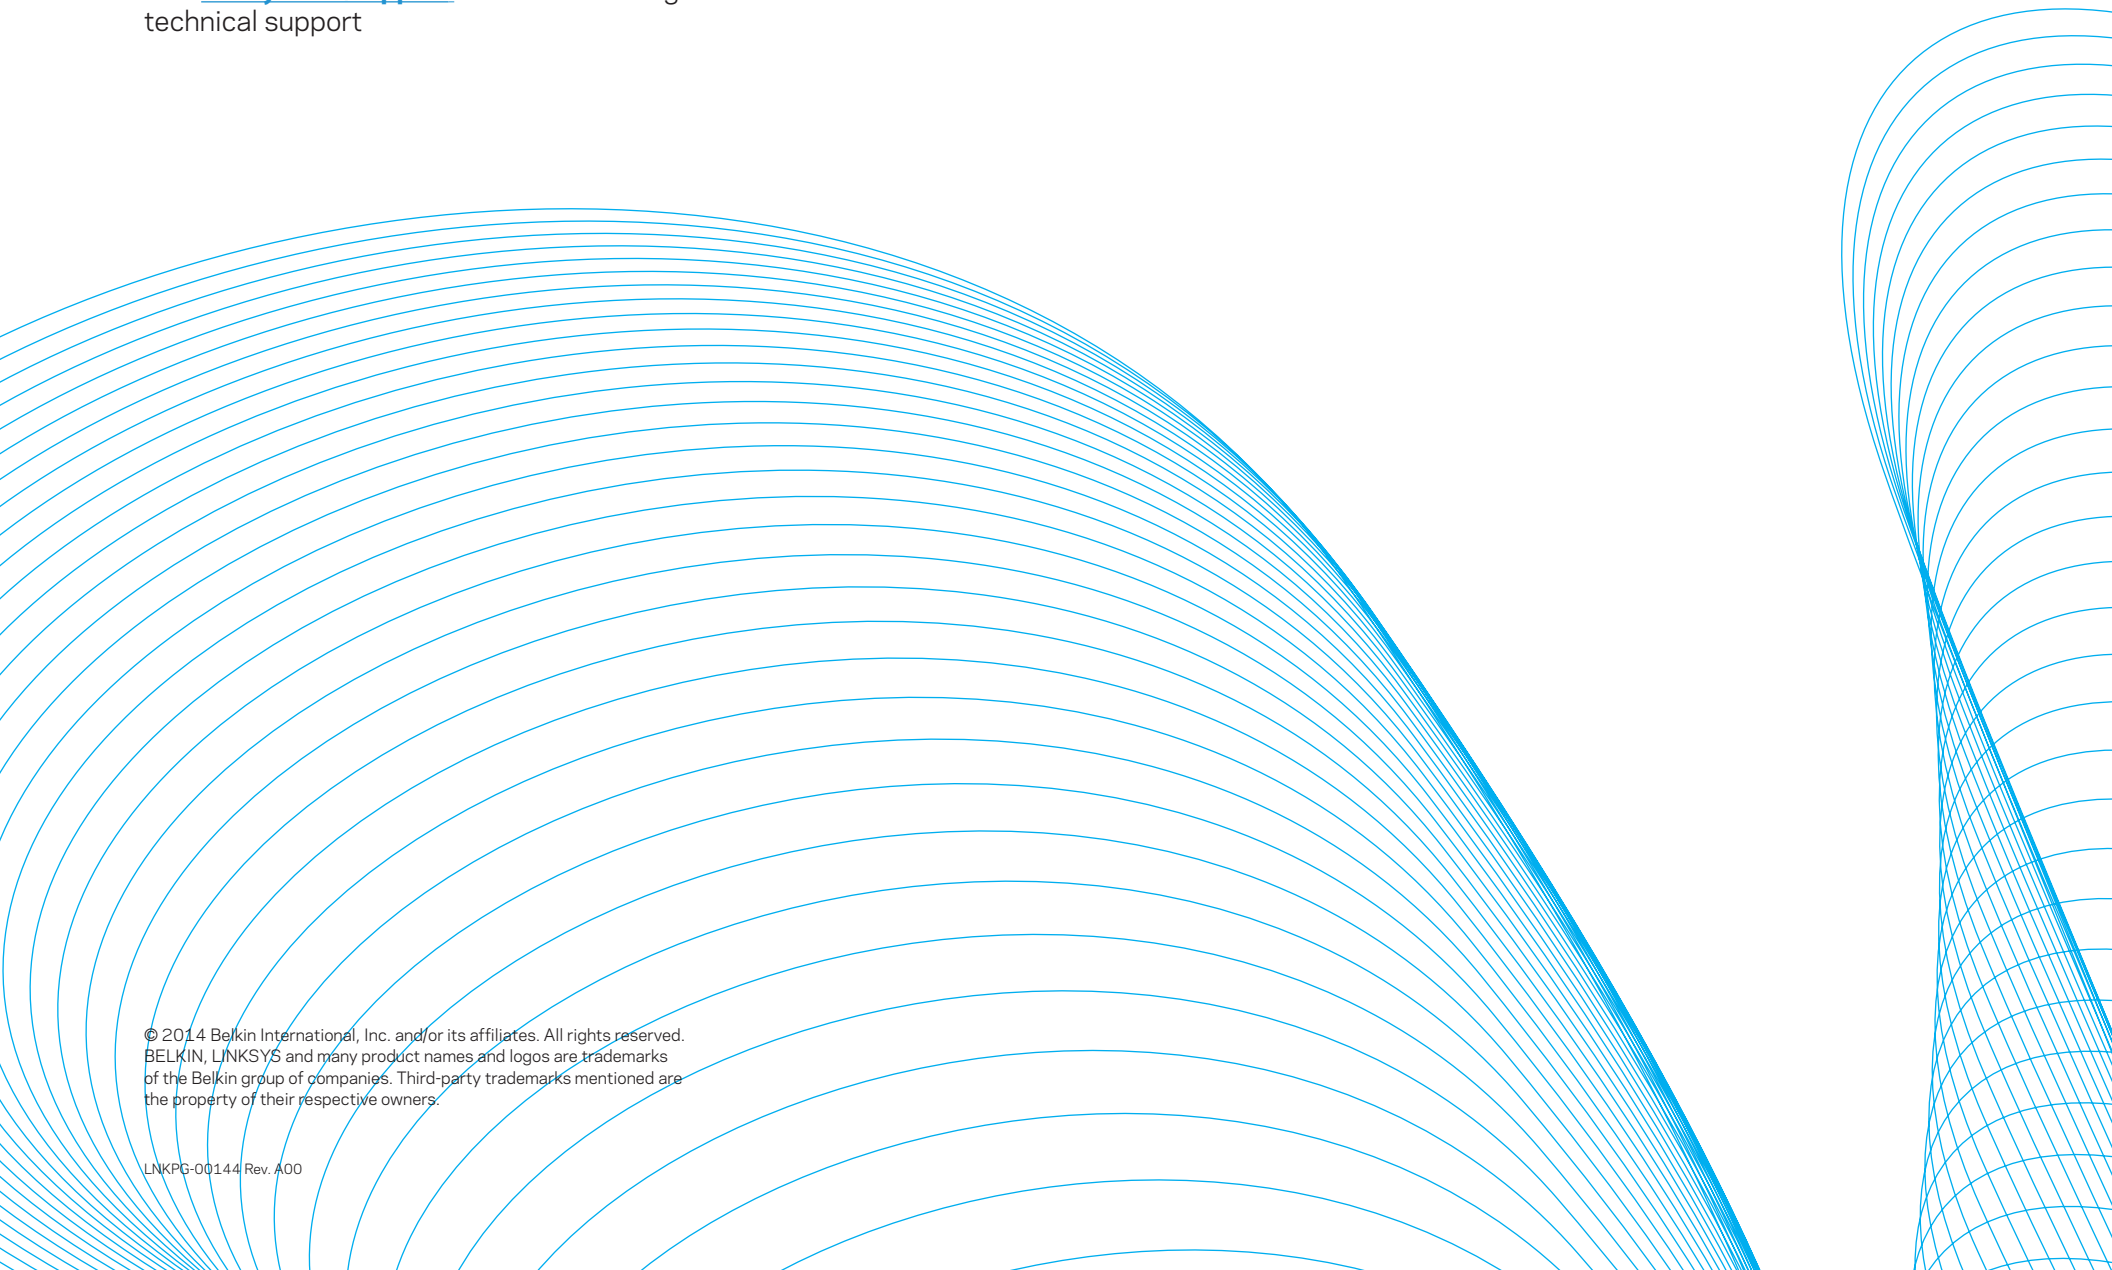

LNKPG-00144 Rev. A00