



User Guide

Managed Switch

LGS5XX

Contents

| | |
|--|------------|
| Chapter 1 - Getting Started | 6 |
| Configuring the Console Port..... | 6 |
| Interface Naming Conventions..... | 8 |
| Window Navigation..... | 9 |
| Chapter 2 - System Status | 11 |
| System Summary..... | 11 |
| RMON..... | 12 |
| Interface Statistics..... | 20 |
| Chapter 3 - Quick Start | 23 |
| Chapter 4 - System Management | 24 |
| System Information..... | 24 |
| TCAM Resources..... | 25 |
| Management Session Timeout | 27 |
| Time..... | 28 |
| SNMP | 36 |
| Logs..... | 58 |
| Chapter 5 - Port Management | 63 |
| Ports..... | 63 |
| Chapter 6 - VLAN Management | 100 |
| Overview..... | 100 |
| VLANs..... | 102 |
| Interfaces..... | 106 |
| VLAN Memberships..... | 110 |
| GVRP..... | 112 |
| VLAN Groups..... | 113 |

| | |
|---|------------|
| Voice VLAN..... | 118 |
| Chapter 7 - Spanning Tree | 123 |
| Overview..... | 123 |
| Spanning Tree..... | 124 |
| STP Interfaces..... | 126 |
| RSTP Interfaces..... | 127 |
| MSTP Properties..... | 130 |
| VLAN to MSTP..... | 130 |
| MSTP Instance Status..... | 133 |
| MSTP Instance Interface..... | 134 |
| Chapter 8 - MAC Address Management | 138 |
| Dynamic MAC Addresses..... | 139 |
| Static MAC Addresses..... | 140 |
| Reserved MAC Addresses | 141 |
| Chapter 9 - Multicast | 143 |
| Overview..... | 143 |
| Feature Configuration | 145 |
| IGMP/MLD Snooping..... | 147 |
| IGMP Snooping..... | 149 |
| MLD Snooping | 151 |
| Multicast Router Ports | 152 |
| Forward All..... | 153 |
| Unregistered Multicast | 155 |
| IGMP/MLD IP Group Addresses | 156 |
| MAC Group Address FDB..... | 158 |
| IP Group Address FDB..... | 161 |
| Chapter 10 - IP Interface | 163 |
| IPv4..... | 163 |
| IPv6..... | 172 |

| | |
|--|------------|
| Chapter 11 - IP Network Operations..... | 185 |
| Domain Name System..... | 185 |
| DHCP..... | 187 |
| IP Source Guard..... | 197 |
| DHCP Snooping Binding Database | 205 |
| ARP Inspection..... | 208 |
| Interface Settings..... | 213 |
| | |
| Chapter 12 - Security | 214 |
| Management Security | 214 |
| RADIUS | 222 |
| Network Access Control..... | 227 |
| Port Security..... | 239 |
| Storm Control | 241 |
| | |
| Chapter 13 - Access Control List | 243 |
| Access Control Lists | 243 |
| MAC-Based ACL/ACE..... | 245 |
| IPv4-Based ACL/ACE | 247 |
| IPv6-Based ACL/ACE | 250 |
| ACL Binding..... | 252 |
| | |
| Chapter 14 - Quality of Service..... | 254 |
| Overview..... | 254 |
| Feature Configuration | 257 |
| Queue Scheduling | 258 |
| CoS/802.1p to Queue | 260 |
| DSCP to Queue..... | 262 |
| Bandwidth Control | 263 |
| Egress Shaping | 265 |
| Basic QoS..... | 266 |
| Advanced QoS..... | 267 |
| QoS Statistics | 279 |

| | |
|-------------------------------|-----|
| Chapter 15 - Maintenance..... | 281 |
| Device Models..... | 281 |
| System Mode & Reboot..... | 282 |
| File Management | 283 |
| Diagnostics..... | 294 |
| Chapter 16 - Support | 301 |
| Appendix | 302 |
| Startup Menu Procedures | 302 |

Chapter 1 - Getting Started

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

- [Configuring with the Console Port](#)
- [Launching the Configuration Utility](#)
- [Interface Naming Conventions](#)
- [Window Navigation](#)

Configuring the Console Port

To configure with the Console Port:

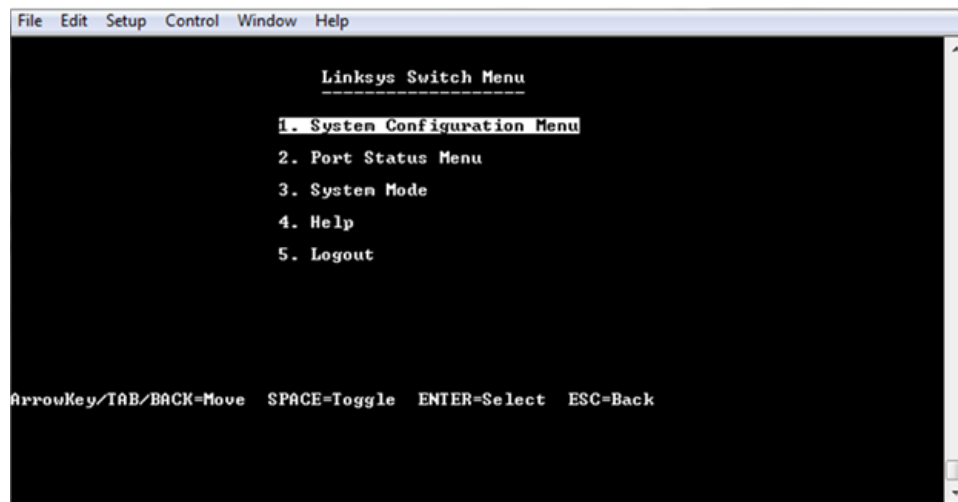
1. Use a provided serial cable to connect to console port
2. Start a terminal application such as Hyper Terminal on your computer
3. Configure the utility with 11520 bit per second, 8 data bits, no parity, 1 stop bit and no flow control. (The firmware supports autobaud detection, the device will detect 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.

Launching the 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)
- Chrome (version 35 and latest)

Note—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.

To open the Web-based configuration utility:

1. Open a Web browser.
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.

To log in:

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

1. Open the GUI. The Login page is displayed.
2. Enter the username/password. The password can contain up to 64 ASCII characters.

To log 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. |
|-------|---|

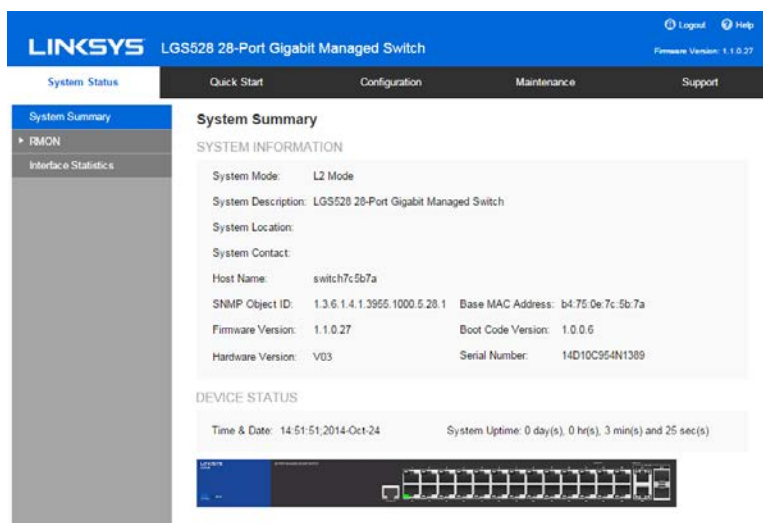
| Button Name | Description |
|------------------|--|
| Close | Click to return to the previous page. Any changes not applied are cleared. |
| Clear All | Click to clear the statistic counters for all interfaces. |
| Clear | Click to clear information, such a counters of an interface or all interface, 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. <ol style="list-style-type: none"> 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

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

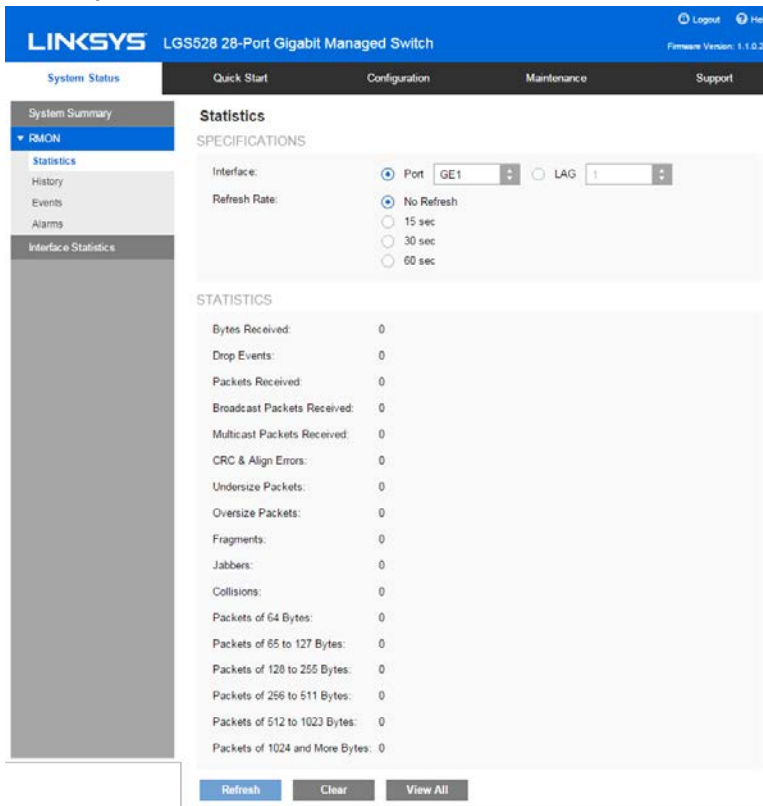
Statistics

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:

1. Click *System Status > RMON > Statistics*.



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

- 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.

RMON History

The RMON feature enables monitoring statistics per interface.

The History Control Table page defines the sampling frequency, amount of samples to store and the port from which to gather the data.

After the data is sampled and stored, it appears in the History Table page that can be viewed by clicking the History button.

To enter RMON control information:

1. Click *System Status > RMON > History*.
2. Click **Add**.

Add History Control

New History Control

History Control Entry Index: 1

History Control Settings

Source Interface: Port GE1 LAG 1

Maximum Samples: 50 (1-50) Sampling Interval: 1800 sec (1-3600)

Owner:

Apply **Close**

3. Enter the parameters.
 - New History Control Entry Index —Displays the number of the new History table entry.
 - Source Interface—Select the type of interface from which the history samples are to be taken.
 - Maximum Samples—Enter the number of samples to store.
 - Samples Collected—RMON is allowed by the standard to not grant all requested samples, but rather to limit the number of samples per request. Therefore, this field represents the sample number actually granted to the request that is equal or less than the requested maximum sample.
 - Sampling Interval—Enter the time in seconds that samples are collected from the ports. The field range is 1-3600.
 - Owner—Enter the RMON station or user that requested the RMON information.
4. Click **Apply**. The entry is added to the History Control Table page, and the Running Configuration file is updated.

5. Click the **History** button (described below) to view the actual statistics.

RMON History Table

The History Table page displays interface-specific statistical network samplings. The samples were configured in the History Control table described above.

To view RMON history statistics:

1. Click *System Status > RMON > History*.
2. Click **History**.
3. From the *History Control Entry Index* drop-down menu, select the entry number of the sample to display.

The fields are displayed for the selected sample.

- Owner—History table entry owner.
- Sample Index—Statistics were taken from this sample.
- Drop Events—Dropped packets due to lack of network resources during the sampling interval. This may not represent the exact number of dropped packets, but rather the number of times dropped packets were detected.
- Bytes Received—Octets received including bad packets and FCS octets, but excluding framing bits.
- Packets Received—Packets received, including bad packets, Multicast, and Broadcast packets.
- Broadcast Packets—Good Broadcast packets excluding Multicast packets.
- Multicast Packets—Good Multicast packets received.
- CRC Align Errors—CRC and Align errors that have occurred.
- Undersize Packets—Undersized packets (less than 64 octets) received.
- Oversize Packets—Oversized packets (over 2000 octets) received.
- Fragments—Fragments (packets with less than 64 octets) received, excluding framing bits, but including FCS octets.
- Jabbers—Total number of received packets that were longer than 2000 octets. This number excludes frame bits, but includes FCS octets that had either a bad FCS (Frame Check Sequence) with an integral number of octets (FCS Error) or a bad FCS with a non-integral octet (Alignment Error) number.
- Collisions—Collisions received.

- Utilization—Percentage of current interface traffic compared to maximum traffic that the interface can handle.

RMON Events

You can control the occurrences that trigger an alarm and the type of notification that occurs.

- Events Page—Configures what happens when an alarm is triggered. This can be any combination of logs and traps.
- Alarms Page—Configures the occurrences that trigger an alarm.

To define RMON events:

1. Click *System Status > RMON > Events*. This page displays previously defined events.
2. Click **Add**.

3. Enter the parameters.
 - Event Entry Index —Displays the event entry index number for the new entry.
 - Community—Enter the SNMP community string to be included when traps are sent (optional). Note that the community must be defined using the Defining SNMPv1,2 Notification Recipients or Defining SNMPv3 Notification Recipients pages for the trap to reach the Network Management Station.
 - Description—Enter a name for the event. This name is used in the Add RMON Alarm page to attach an alarm to an event.
 - Notification Type—Select the type of action that results from this event.
 - None—No action occurs when the alarm goes off.
 - Event Log (Event Log Table)—Add a log entry to the Event Log table when the alarm is triggered.

- Trap (SNMP Manager and SYSLOG Server)—Send a trap to the remote log server when the alarm goes off.
 - Trap and Event Log—Add a log entry to the Event Log table and send a trap to the remote log server when the alarm goes off.
 - Last Event Time—Displays the time of the event. (This is a read-only table in the parent window and cannot be defined).
 - Owner—Enter the device or user that defined the event.
4. Click **Apply**. The RMON event is saved to the Running Configuration file.
 5. Click **Event Log** to display the log of alarms that have occurred and that have been logged (see description below).

RMON Events Logs

The Event Log Table page displays the log of events (actions) that occurred. Two types of events can be logged: Log or Log and Trap. The action in the event is performed when the event is bound to an alarm (see the Alarms page) and the conditions of the alarm have occurred.

1. Click *System Status > RMON > Events*.
2. Click Event Log.

This page displays the following fields:

- Event Entry Index —Event's log entry number.
- Log No.—Log number (within the event).
- Log Time—Time that the log entry was entered.
- Description—Description of event that triggered the alarm.

RMON Alarms

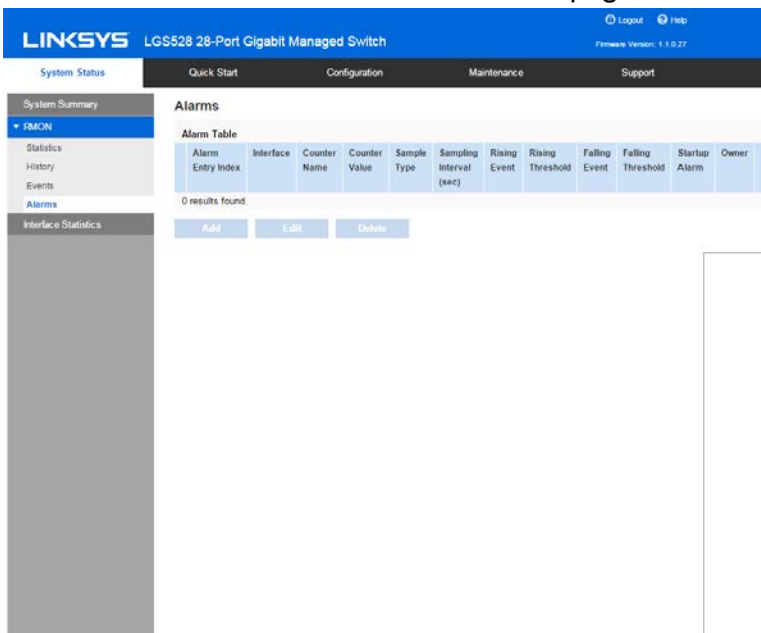
RMON alarms provide a mechanism for setting thresholds and sampling intervals to generate exception events on counters or any other SNMP object counter maintained by the agent. Both the rising and falling thresholds must be configured in the alarm. After a rising threshold is crossed, no rising events are generated until the companion falling threshold is crossed. After a falling alarm is issued, the next alarm is issued when a rising threshold is crossed.

One or more alarms are bound to an event, which indicates the action to be taken when the alarm occurs.

Alarm counters can be monitored by either absolute values or changes (delta) in the counter values.

To enter RMON alarms:

1. Click *System Status > RMON > Alarms*. All previously-defined alarms are displayed. The fields are described in the Add RMON Alarm page below.



In addition to those fields, the following field appears:

- Counter Value—Displays the value of the statistic during the last sampling period.
2. Click **Add**.
 3. Enter the parameters.
 - Alarm Entry Index—Displays the alarm entry number.
 - Interface—Select the type of interface for which RMON statistics are displayed.
 - Counter Name—Select the MIB variable that indicates the type of occurrence measured.
 - Sample Type—Select the sampling method to generate an alarm.

The options:

- Absolute—If the threshold is crossed, an alarm is generated.
 - Delta—Subtracts the last sampled value from the current value. The difference in the values is compared to the threshold. If the threshold was crossed, an alarm is generated.
- Interval—Enter the alarm interval time in seconds.

- Startup Alarm—Select the first event from which to start generation of alarms. Rising is defined by crossing the threshold from a low-value threshold to a higher-value threshold.
 - Rising Alarm—A rising value triggers the rising threshold alarm.
 - Falling Alarm—A falling value triggers the falling threshold alarm.
 - Rising and Falling Alarm—Both rising and falling values trigger the alarm.
 - Owner—Enter the name of the user or network management system that receives the alarm.
 - Rising Threshold—Enter the value that triggers the rising threshold alarm.
 - Rising Event—Select an event to be performed when a rising event is triggered. Events are created in the Events page.
 - Falling Threshold—Enter the value that triggers the falling threshold alarm.
 - Falling Event—Select an event to be performed when a falling event is triggered.
4. Click Apply. The RMON alarm is saved to the Running Configuration file.

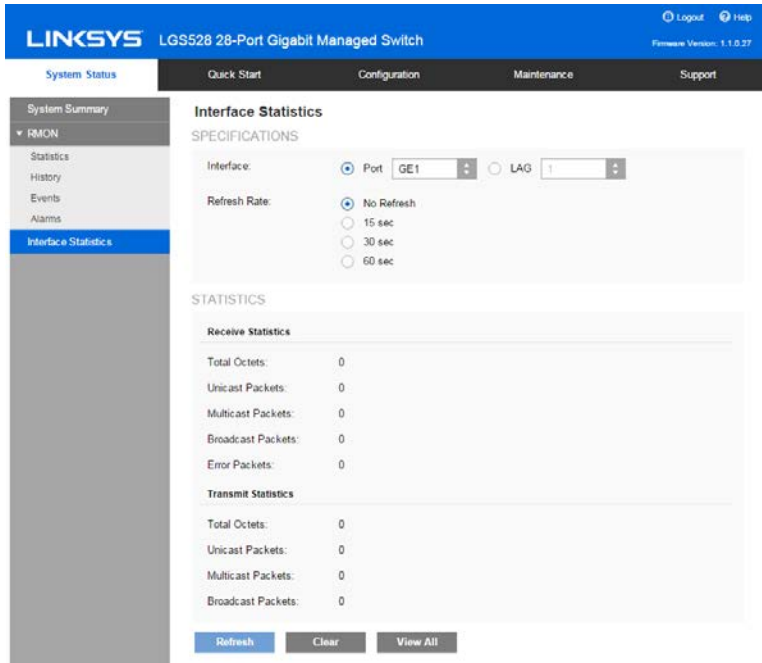
Interface Statistics

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:

1. Click *System Status > Interface Statistics*.



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

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

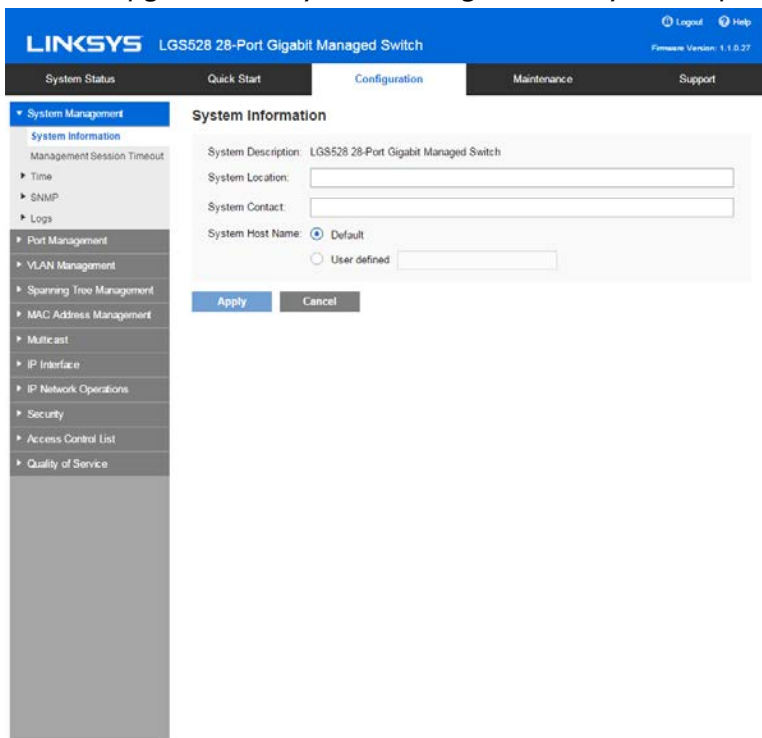
This chapter describes the following topics:

- [System Information](#)
- [TCAM Resources](#)
- [Management Session Timeout](#)
- [Time](#)
- [SNMP](#)
- [Logs](#)

System Information

To enter system information:

1. Click *Configuration > System Management > System Information*.



The screenshot displays the Linksys web interface for an LGS528 28-Port Gigabit Managed Switch. The top navigation bar includes 'System Status', 'Quick Start', 'Configuration', 'Maintenance', and 'Support'. The 'Configuration' tab is active, and the left sidebar shows 'System Management' expanded to 'System Information'. The main content area is titled 'System Information' and contains the following fields:

- System Description: LGS528 28-Port Gigabit Managed Switch
- System Location: [Text input field]
- System Contact: [Text input field]
- System Host Name: Default, User defined [Text input field]

At the bottom of the form are 'Apply' and 'Cancel' buttons.

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 is 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).
3. Click **Apply** to save the values in the Running Configuration file.

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:

1. Click *Configuration > System Management > TCAM Resources*.

LINKSYS LGS528 28-Port Gigabit Managed Switch Firmware Version: 1.1.0.27

System Status Quick Start **Configuration** Maintenance Support

System Management

- System Information
- TCAM Resources**
- Management Session Timeout
- Time
- SNMP
- Logs
- Port Management
- VLAN Management
- Spanning Tree Management
- MAC Address Management
- Multicast
- IP Interface
- IP Network Operations
- Security
- Access Control List
- Quality of Service

TCAM Resources

IPv4 TCAM SPECIFICATION

Maximum IPv4 TCAM Entries: Use Default User Defined (8-476)

IPv4 TCAM USAGE

| | Count | TCAM Entries |
|------------------|-------|--------------|
| IPv4 Hosts: | 2 | 2 |
| IPv4 Interfaces: | 1 | 2 |
| IPv4 Routes: | 1 | 1 |
| Total: | | 5 |

OTHER TCAM USAGE

| | In-Use | Maximum Allocated |
|---------|--------|-------------------|
| Non-IP: | 10 | 361 |

Apply Cancel

2. 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.

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
 - Count—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.

Counters are displayed for Non-IP TCAM Usage:

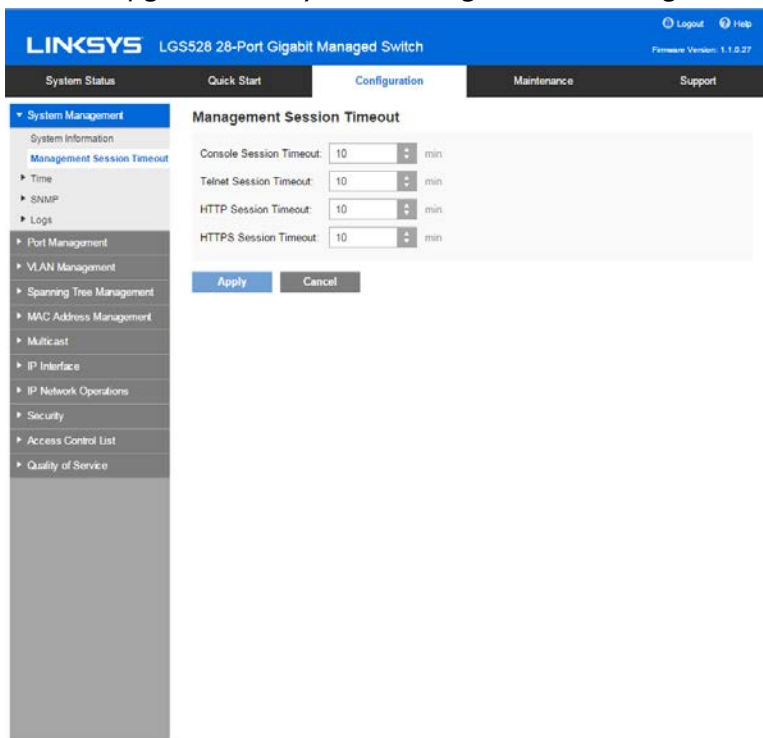
- Non-IP
 - In Use—Number of TCAM entries currently used by applications and features, excluding IP routing.
 - Maximum Allocated—Number of available TCAM entries that can be used by applications and features, excluding IP routing.

Management Session Timeout

The Management Session Timeout configures the time intervals that the management sessions can remain idle before they timeout and you must log in again to reestablish the session.

To set the idle session timeout for various types of sessions:

1. Click *Configuration > System Management > Management Session Timeout*.



2. Select the timeout for the following sessions from the corresponding list. The default timeout value is 10 minutes.
 - Console Session Timeout—Select the timeout for a console session.
 - Telnet Session Timeout—Select the timeout for a Telnet session.

- HTTP Session Timeout—Select the timeout for an HTTP session.
 - HTTPS Session Timeout—Select the timeout for an HTTPS session.
3. Click **Apply** to set the configuration settings on the device.

Time

This section describes the options for configuring the system time, time zone, and Daylight Savings Time (DST).

- [Overview](#)
- [System Time](#)
- [SNTP Unicast Server](#)
- [SNTP Multicast/Anycast](#)

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

System time can be set manually by the user, dynamically from an SNTP server, or synchronized from the PC running the GUI. 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 the PC running the GUI, SNTP, values set manually, or if all else fails, from the factory defaults.

Time

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

- Manual—User 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 hostname, 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, where:
 - 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 system time from an SNTP server in one of the following ways:

- Client Broadcast Reception (passive mode)—SNTP servers broadcast the time, and the device listens to these broadcasts. When the device is in this mode, there is no need to define a Unicast SNTP server.
- Client Broadcast 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:

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

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

The screenshot shows the 'System Time' configuration page in the Linksys web interface. The page title is 'LINKSYS LGS528 28-Port Gigabit Managed Switch' with 'Firmware Version: 1.1.0.27'. The navigation menu includes 'System Status', 'Quick Start', 'Configuration', 'Maintenance', and 'Support'. The left sidebar shows 'System Management' expanded, with 'Time' selected. The main content area is titled 'System Time' and contains the following sections:

- Current Time:** 15:07:25 2014-Oct-24
- Clock Source:** Includes checkboxes for SNTP, SNTP IPv4 Multicast Rx, SNTP IPv6 Multicast Rx, SNTP Client Unicast, SNTP IPv4 Anycast Tx, and SNTP IPv6 Anycast Tx.
- Manual Date/Time:** 2014-Oct-24 yyyy-mm-dd 15:07 hh:mm
- Time Zone:** Includes 'Time Zone from DHCP' (disabled), 'DHCP Time Zone' (N/A), 'Time Zone Offset' (UTC), and 'Time Zone Acronym'.
- Daylight Saving Time:** Includes 'Daylight Saving' (disabled), 'Time Set Offset (1-1440): 60 min', 'Daylight Savings Type' (USA selected), and 'From'/'To' date and time fields.

Buttons for 'Apply' and 'Cancel' are at the bottom.

2. Enter the following 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 from any 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 from any 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 from 1-1440. 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. This allows customization of the start and stop of DST.
 - Recurring - DST occurs on the same date every year. This allows customization of the start and stop of DST.

- For Daylight Savings Time, enter the following parameters:
 - From - Date and time that DST starts.
 - To - Date and time that DST ends.
- Recurring From, enter the following parameters that indicate 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.
- (Recurring) To - Enter the following parameters that indicate when DST ends each year:
 - 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.

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 (in the [System Time](#) page).

To view the Unicast SNTP server page:

- Click *Configuration > System Management > Time > SNTP Unicast Server*.

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

1. SNTP Server—SNTP server IP address. The preferred server, or hostname, is chosen according to its stratum level.
2. SNTP Server Status—SNTP server status. The possible values:
 - 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 does not fully trust its own time server (i.e. when first booting up the SNTP server).
3. 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.
 4. Offset—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.
 5. Delay—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.
 6. Poll Interval—Displays whether polling is enabled or disabled.
 7. Last Response Time—Last date and time a response was received from this SNTP server.

To add a Unicast SNTP server, enable SNTP Client Unicast.

1. Click **Add**.

2. 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 6 or Version 4.

- IPv6 Address Type—Select the IPv6 address type (if IPv6 is used). The options:
 - 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.
 - Link Local Interface—Select the link local interface (if IPv6 Address Type Link Local is selected) from the list.Global—The IPv6 address is a global Unicast IPV6 type that is visible and reachable from other networks.
- 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 the name of an 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.

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

SNTP Multicast/Anycast

This page is only available in Layer 3.

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

1. Click *Configuration > System Management > Time > SNTP Multicast/Anycast*.
2. Click **Add** to select the interface for SNTP reception/transmission.

Add SNTP Multicast/Anycast Interface

Enter New Interface

Interface: Port LAG VLAN

Select an interface.

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. It covers the following topics:

- [SNMP Versions and Workflow](#)
- [Feature Configuration](#)
- [Views](#)
- [Groups](#)
- [Users](#)
- [Communities](#)
- [Notification Filters](#)
- [V1/V2 Notification Recipients](#)
- [V3 Notification Recipients](#)

SNMP Versions and Workflow

The device functions as SNMP agent and supports SNMPv1, v2, and v3. It also reports system events to trap receivers using the traps defined in the supported MIBs (Management Information Base).

SNMPv1 and v2

To control access to the system, a list of community entries is defined. Each community entry consists of a community string and its access privilege. The system responds only to SNMP messages specifying the community which has the correct permissions and correct operation.

SNMP agents maintain a list of variables that are used to manage the device. These variables are defined in the Management Information Base (MIB).

Note—Due to the security vulnerabilities of other versions, it is recommended to use SNMPv3.

SNMPv3

In addition to the functionality provided by SNMPv1 and v2, SNMPv3 applies access control and new trap mechanisms to SNMPv1 and SNMPv2 PDUs. SNMPv3 also defines a User Security Model (USM) that includes:

- Authentication—Provides data integrity and data origin authentication.
- Privacy—Protects against disclosure message content. Cipher Block- Chaining (CBC-DES) is used for encryption. Either authentication alone can be enabled on an SNMP message, or both authentication and privacy can be enabled on an SNMP message. However, privacy cannot be enabled without authentication.
- Timeliness—Protects against message delay or playback attacks. The SNMP agent compares the incoming message time stamp to the message arrival time.

SNMP Workflow

Note—For security reasons, SNMP is disabled by default. Before you can manage the device via SNMP, you must turn on SNMP in the *SNMP>Feature Configuration* page.

The following is the recommended series of actions for configuring SNMP:

If you decide to use SNMPv1 or v2:

1. Click *Configuration > System Management >SNMP > Communities*.
2. Click **Add**.

Add SNMP Community

Enter New Community

SNMP Management Station: All User Defined

IP Version: IPv4 IPv6

IPv6 Address Type: Global Link Local Interface: VLAN 1

IP Address:

Community:

Community Setting

Access Control: Basic Advanced

Access Mode: Read Only Read Write SNMP Admin

View Name: Default

Group Name:

The community can be associated with access rights and a view in Basic mode or with a group in Advanced mode. There are two ways to define access rights of a community:

- Basic mode—The access rights of a community can be configured with Read Only, Read Write, or SNMP Admin. In addition, you can restrict the access to the community to only certain MIB objects by selecting a view (defined in the Views page).
 - Advanced Mode—The access rights of a community are defined by a group (defined in the Groups page). You can configure the group with a specific security model. The access rights of a group are Read, Write, and Notify.
3. Choose whether to restrict the SNMP management station to one address or allow SNMP management from all addresses. If you choose to restrict SNMP management to one address, then input the address of your SNMP Management PC in the IP Address field.
 4. Input the unique community string in the Community String field.
 5. Optionally, enable traps by using the Trap Settings page.
 6. Optionally, define a notification filter(s) by using the Notification Filter page.
 7. Configure the notification recipients on the Notification Recipients SNMPv1, 2 page.

If you decide to use SNMPv3:

1. Define the SNMP engine by using the Engine ID page. Either create a unique Engine ID or use the default Engine ID. Applying an Engine ID configuration clears the SNMP database.
2. Optionally, define SNMP view(s) by using the Views page. This limits the range of OIDs available to a community or group.
3. Define groups by using the Groups page.
4. Define users by using the SNMP Users page, where they can be associated with a group. If the SNMP Engine ID is not set, then users may not be created.
5. Optionally, enable or disable traps by using the Trap Settings page.
6. Optionally, define a notification filter(s) by using the Notification Filter page.
7. Define a notification recipient(s) by using the Notification Recipients SNMPv3 page.

Model OIDs

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

| Mode Name | Description | Object ID |
|-----------|---|--|
| LGS528P | 26-ports Gigabit PoE+ Management Switch + 2 COMBO ports | enterprises(1) .linksys(3955) .smb(1000) .5 28 .2 |
| LGS528 | 26-ports Gigabit Management Switch + 2 COMBO ports | enterprises(1) .linksys(3955) .smb(1000) .5 .28 .1 |
| LGS552P | 50-ports Gigabit PoE+ Management Switch + 2 COMBO ports | enterprises(1) .linksys(3955) .smb(1000) .5 .52 .2 |
| LGS552 | 50-ports Gigabit Management Switch + 2 COMBO ports | enterprises(1) .linksys(3955) .smb(1000) .5 .52 .1 |

Private OIDs are placed under: enterprises(1).linksys(3955).smb(1000).switch01(201).

Feature Configuration

The Engine ID is used by SNMPv3 entities to uniquely identify them. An SNMP agent is considered an authoritative SNMP engine. This means that the agent responds to incoming messages (Get, GetNext, GetBulk, Set) and sends trap messages to a manager. The agent's local information is encapsulated in fields in the message.

Each SNMP agent maintains local information that is used in SNMPv3 message exchanges. The default SNMP Engine ID is comprised of the enterprise number and the default MAC address. This engine ID must be unique for the administrative domain, so that no two devices in a network have the same engine ID.

Local information is stored in four MIB variables that are read-only (snmpEngineId, snmpEngineBoots, snmpEngineTime, and snmpEngineMaxMessageSize).

Caution—When the engine ID is changed, all configured users and groups are erased.

To configure SNMP:

1. Click *Configuration > System Management > SNMP > Feature Configuration*.

The screenshot shows the configuration page for a Linksys LGS528 28-Port Gigabit Managed Switch. The page title is "LINKSYS LGS528 28-Port Gigabit Managed Switch" with a firmware version of 1.1.0.27. The navigation menu includes System Status, Quick Start, Configuration, Maintenance, and Support. The left sidebar shows a tree view with System Management expanded, containing System Information, TCAM Resources, Management Session Timeout, Time, and SNMP. Under SNMP, Feature Configuration is selected. The main content area is titled "Feature Configuration" and contains the following settings:

- SNMP: Enable
- SNMP Notification: Enable
- Authentication Notification: Enable
- Local SNMPv3 Engine ID: Use Default, None, User Defined (with a text input field and "(10-64 hex digits)" label)

Below the settings are "Apply" and "Cancel" buttons. Underneath is a table titled "SNMPv3 Remote Engine Table" with columns "Engine IP Address" and "Engine ID". The table is currently empty, showing "0 results found." Below the table are "Add", "Edit", and "Delete" buttons.

2. Enter the following fields:

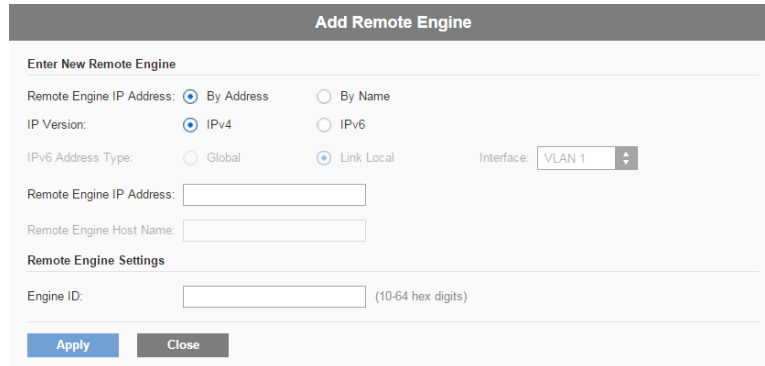
- SNMP—Select to enable SNMP.
- Authentication Notification—Select to enable SNMP authentication failure notification.
- SNMP Notification—Select to enable SNMP notifications.
- Local SNMPv3 Engine ID—Configure the engine. The options are:
 - Use Default—Select to use the device-generated engine ID. The default engine ID is based on the device MAC address, and is defined per standard as:
 - First 4 octets—First bit = 1, the rest is the IANA enterprise number.
 - Fifth octet—Set to 3 to indicate the MAC address that follows.
 - Last 6 octets—MAC address of the device.
 - None—No engine ID is used.
 - User Defined—Enter the local device engine ID. The field value is a hexadecimal string (range: 10 - 64). Each byte in the hexadecimal character strings is represented by two hexadecimal digits.

All remote engine IDs and their IP addresses are displayed in the Remote Engine ID table.

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

The Remote Engine ID table shows the mapping between IP addresses of the engine and Engine ID. To add the IP address of an engine ID:

4. Click **Add**.



The screenshot shows a configuration window titled "Add Remote Engine". It contains the following fields and options:

- Enter New Remote Engine** (Section Header)
- Remote Engine IP Address:** Radio buttons for By Address and By Name.
- IP Version:** Radio buttons for IPv4 and IPv6.
- IPv6 Address Type:** Radio buttons for Global and Link Local. An **Interface:** dropdown menu is set to "VLAN 1".
- Remote Engine IP Address:** A text input field.
- Remote Engine Host Name:** A text input field.
- Remote Engine Settings** (Section Header)
- Engine ID:** A text input field with a note "(10-64 hex digits)".
- Buttons for **Apply** and **Close**.

Enter the following fields:

- Remote Engine IP Address—Select whether to specify the Engine ID server by IP address or name.
 - IP Version—Select the supported IP format.
 - 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.
 - Remote Engine IP Address—Enter the IP address of the log server.
 - Remote Engine IP Name—Enter the domain name of the log server.
 - Engine ID—Enter the Engine ID.
5. Click **Apply**. The Running Configuration file is updated.

Views

A view is a user-defined label for a collection of MIB subtrees. Each subtree ID is defined by the Object ID (OID) of the root of the relevant subtrees. Either well-known names can be used to specify the root of the desired subtree or an OID can be entered (see Model OIDs).

Each subtree is either included or excluded in the view being defined.

The Views page enables creating and editing SNMP views. The default views (Default, DefaultSuper) cannot be changed.

Views can be attached to groups in the Groups page or to a community which employs basic access mode through the Communities page.

To define SNMP views:

1. Click *Configuration > System Management > SNMP > Views*.

| View Name | Object ID | Object View |
|--------------|---------------------------------|-------------|
| Default | 1 | Included |
| Default | 1.3.6.1.6.3.13 | Excluded |
| Default | 1.3.6.1.6.3.16 | Excluded |
| Default | 1.3.6.1.6.3.18 | Excluded |
| Default | 1.3.6.1.6.3.12.1.2 | Excluded |
| Default | 1.3.6.1.6.3.12.1.3 | Excluded |
| Default | 1.3.6.1.6.3.16.1.2 | Excluded |
| Default | 1.3.6.1.4.1.3965.1000.201.98.1 | Excluded |
| Default | 1.3.6.1.4.1.3965.1000.201.2.7.2 | Excluded |
| DefaultSuper | 1 | Included |

2. Click **Add** to define new views.

3. Enter the parameters.

- View Name—Enter a view name between 0-30 characters)
- View Object—Select the node in the MIB tree that is included or excluded in the selected SNMP view.

The options to select the object:

- Object ID—Enter an OID not offered in the Object ID Selection List option.
- Object ID Selection List—Enables you to navigate the MIB tree. Press the Up arrow to go to the level of the selected node's parent and siblings; press the Down arrow to descend to the level of the selected node's children. Click nodes in the view to pass from one node to its sibling. Use the scrollbar to bring siblings in view.

4. Include or exclude the MIB object from the view. If Include Object is selected, the MIB objects are included in the view, otherwise they are excluded.

5. Click **Apply**.

6. In order to verify your view configuration, select the user-defined views from the Filter: View Name list.

The following views exist by default:

- Default—Default SNMP view for read and read/write views.
- DefaultSuper—Default SNMP view for administrator views. Other views can be added.
- Object ID—Displays the Object ID and its subtree to be included or excluded in the SNMP view.
- Object View—Displays whether the defined object and its subtree are included or excluded in the selected SNMP view.

Groups

In SNMPv1 and SNMPv2, a community string is sent along with the SNMP frames. The community string acts as a password to gain access to an SNMP agent. However, neither the frames nor the community string are encrypted. Therefore, SNMPv1 and SNMPv2 are not secure.

In SNMPv3, the following security mechanisms can be configured.

- Authentication—The device checks that the SNMP user is an authorized system administrator. This is done for each frame.
- Privacy—SNMP frames can carry encrypted data. Thus, in SNMPv3, there are three levels of security:
 - No security (No authentication and no privacy)
 - Authentication (Authentication and no privacy)
 - Authentication and privacy

SNMPv3 provides a means of controlling the content each user can read or write and the notifications they receive. A group defines read/write privileges and a level of security. It becomes operational when it is associated with an SNMP user or community.

Note—*To associate a non-default view with a group, first create the view in the Views page.*

To create an SNMP group:

1. Click *Configuration > System Management > SNMP > Groups*.

This page displays the existing SNMP groups and their security levels. The following fields are displayed for each SNMP group (only the fields not explained in the Add page):

- No Authentication Read View—No authentication is needed, and anyone is able to read the view.
- No Authentication Write View—No authentication is needed, and anyone is able to write the view.
- No Authentication Notify View—No authentication is needed, and anyone is able to receive notification of the view.
- Authentication Read View—Only authenticated users are allowed to read the view. By default, all users or community of a group can access all the MIB objects. A group can be limited to specific view(s) based on the read, write, notify authentication and/or privacy configurations.
- Authentication Write View—Only authenticated users are able to write the view. Management access is write for the selected view.

- Authentication Notify View—Only authentication users are allowed to received notification.
- Privacy Read View— When reading the objects in the view, the SNMP messages are encrypted.
- Privacy Write View—When writing the object in the view, the SNMP messages are encrypted.
- Privacy Notify View - Notification on the objects in the view are encrypted.

2. Click **Add**.

3. Enter the parameters.

- Group Name—Enter a new group name.
- Security Model—Select the SNMP version attached to the group, SNMPv1, v2, or v3.

Three types of views with various security levels can be defined. For each security level, select the views for Read, Write and Notify by entering the following fields:

- Enable—Select this field to enable the Security Level.
- Security Level—Define the security level attached to the group. SNMPv1 and SNMPv2 support neither authentication nor privacy. If SNMPv3 is selected, select to enable one of the following:
 - o No Authentication and No Privacy—Neither the Authentication nor the
 - o Privacy security levels are assigned to the group.
 - o Authorized View—Select the Read, Write and Notify views associated with this group and with the above security level.

- o Authentication and No Privacy—Authenticates SNMP messages, and ensures the SNMP message origin is authenticated but does not encrypt them.
- o Authorized View—Select the Read, Write and Notify views associated with this group and with the above security level.
- o Authentication and Privacy—Authenticates SNMP messages, and encrypts them.
- o Authorized View—Select the Read, Write and Notify views associated with this group and with the above security level.

4. Click **Apply**. The SNMP group is saved to the Running Configuration file.

Users

The screenshot shows the Linksys configuration interface for an LGS528 28-Port Gigabit Managed Switch. The page is titled "Users" and displays a "User Table" with the following columns: User Name, Engine IP Address, Engine ID, Group Name, Authentication Method, and Privacy Method. Below the table, it indicates "0 results found." and provides "Add", "Edit", and "Delete" buttons. The left sidebar shows the navigation menu with "System Management" expanded to "SNMP" and "Users" selected. The top navigation bar includes "System Status", "Quick Start", "Configuration", "Maintenance", and "Support".

An SNMP user is defined by the login credentials (username, passwords, and authentication method) and by the context and scope in which it operates by association with a group and an Engine ID.

The configured user has the attributes of its group, having the access privileges configured within the associated view.

Groups enable network managers to assign access rights to a group of users instead of to a single user.

A user can only belong to a single group.

To create an SNMPv3 user, the following must first exist:

- An engine ID must first be configured on the device. This is done in the Engine ID page.
- An SNMPv3 group must be available. An SNMPv3 group is defined in the Groups page.

To display SNMP users and define new ones:

1. Click *Configuration > System Management > SNMP > Users*.

This page contains existing users.

2. Click **Add**.

This page provides information for assigning SNMP access control privileges to SNMP users.

The screenshot shows the 'Add User' configuration page. At the top is a dark grey header with the text 'Add User'. Below the header is a section titled 'Enter New User' containing a text input field for 'User Name:'. Underneath is the 'Engine ID:' section with two radio buttons: 'Local' (selected) and 'Engine' (with a dropdown arrow). Below that is the 'User Settings' section, which includes a dropdown menu for 'Group Name:', three radio buttons for 'Authentication Method:' (None, selected; MD5; SHA), a text input for 'Authentication Password:', two radio buttons for 'Privacy Method:' (None, selected; DES), and another text input for 'Privacy Password:'. At the bottom of the form are two buttons: 'Apply' (blue) and 'Close' (grey).

3. Enter the parameters.

- **User Name**—Enter a name for the user.
- **Engine ID**—Select either the local or remote SNMP entity to which the user is connected. Changing or removing the local SNMP Engine ID deletes the SNMPv3 User Database. To receive inform messages and request information, you must define both a local and remote user.
 - **Local**—User is connected to the local device.
 - **Engine**—User is connected to a different SNMP entity besides the local device. If the remote Engine ID is defined, remote devices receive inform messages, but cannot make requests for information.
- **Group Name**—Select the SNMP group to which the SNMP user belongs. SNMP groups are defined in the Add Group page.

Note—Users who belong to groups which have been deleted remain, but they are inactive.

- Authentication Method—Select the Authentication method that varies according to the Group Name assigned. If the group does not require authentication, then the user cannot configure any authentication. The options are:
 - None—No user authentication is used.
 - MD5—A password that is used for generating a key by the MD5 authentication method.
 - SHA—A password that is used for generating a key by the SHA (Secure Hash Algorithm) authentication method.
- Authentication Password—If authentication is accomplished by either a MD5 or a SHA password, enter the local user password in either Encrypted or Plaintext. Local user passwords are compared to the local database, and can contain up to 32 ASCII characters.
- Privacy Method—Select one of the following options:
 - None—Privacy password is not encrypted.
 - DES—Privacy password is encrypted according to the Data Encryption Standard (DES).
- Privacy Password—16 bytes are required (DES encryption key) if the DES privacy method was selected. This field must be exactly 32 hexadecimal characters. The Encrypted or Plaintext mode can be selected.

4. Click **Apply** to save the settings.

Communities

Access rights in SNMPv1 and SNMPv2 are managed by defining communities in the *Communities* page. The community name is a type of shared password between the SNMP management station and the device. It is used to authenticate the SNMP management station.

Communities are only defined in SNMPv1 and SNMPv2 because SNMPv3 works with users instead of communities. The users belong to groups that have access rights assigned to them.

The Communities page associates communities with access rights, either directly (Basic mode) or through groups (Advanced mode):

- Basic mode—The access rights of a community can configure with Read Only, Read Write, or SNMP Admin. In addition, you can restrict the access to the community to only certain MIB objects by selecting a view (defined in the SNMP Views page).

- **Advanced Mode**—The access rights of a community are defined by a group (defined in the Groups page). You can configure the group with a specific security model. The access rights of a group are Read, Write, and Notify.

To define SNMP communities:

1. Click *Configuration > System Management > SNMP > Communities*.

This page contains a table of configured SNMP communities and their properties.

2. Click **Add**.

This page enables network managers to define and configure new SNMP communities.

3. Enter the following fields:

- **SNMP Management Station**—Click User Defined to enter the management station IP address that can access the SNMP community. Click All to indicate that any IP device can access the SNMP community.
- **IP Version**—Select either IPv4 or IPv6.
- **IPv6 Address Type**—Select the supported 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**—If the IPv6 address type is Link Local, select whether it is received through a VLAN or ISATAP.
- **IP Address**—Enter the SNMP management station IP address.

- **Community**—Enter the community name used to authenticate the management station to the device.
- **Access Control**—Select one of the following:
 - **Basic**—In this mode, there is no connection to any group. You can only choose the community access level (Read Only, Read Write, or SNMP Admin) and, optionally, further qualify it for a specific view. By default, it applies to the entire MIB.
 - **Advanced**—In this mode, access is controlled by group configurations.
- **Access Mode**—Configure the community:
 - **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.
 - **SNMP Admin**—User has access to all device configuration options, as well as permissions to modify the community. SNMP Admin is equivalent to Read Write for all MIBs except for the SNMP MIBs. SNMP Admin is required for access to the SNMP MIBs.
 - **View Name**—Select an SNMP view (a collection of MIB subtrees to which access is granted).
 - **Group Name**—Select an SNMP group that determines the access rights in Advanced mode.
- Click **Apply**. The SNMP Community is defined, and the Running Configuration is updated.

Notification Filters

The screenshot shows the Linksys web interface for the LGS528 28-Port Gigabit Managed Switch. The top navigation bar includes 'System Status', 'Quick Start', 'Configuration' (selected), 'Maintenance', and 'Support'. The left sidebar lists various system management options, with 'SNMP' expanded to show 'Notification Filters' as the active selection. The main content area is titled 'Notification Filters' and contains a 'Notification Filters Table' with a search bar and a table with columns for 'Filter Name', 'Object ID', and 'Object Filter'. The table currently shows '0 results found.' and has 'Add', 'Edit', and 'Delete' buttons below it. The footer contains the copyright notice: '© 2014 Belkin International, Inc. and/or its subsidiaries and affiliates, including Linksys, LLC. All rights reserved.'

The Notification Filter page enables configuring SNMP notification filters and Object IDs (OIDs) that are checked. After creating a notification filter, it is possible to attach it to a notification recipient in the Notification Recipients SNMPv1/v2 page, and Notification Recipients SNMPv3 page.

The notification filter enables filtering the type of SNMP notifications that are sent to the management station based on the OID of the notification to be sent.

To define a notification filter:

1. Click *Configuration > System Management>SNMP > Notification Filter*.

The Notification Filter page contains notification information for each filter. The table is able to filter notification entries by Filter Name.

2. Click **Add**.
3. Enter the parameters.

The screenshot shows the 'Add Notification Filter' configuration page. At the top, there is a dark grey header with the text 'Add Notification Filter'. Below the header, the page is divided into two main sections: 'Enter New Filter' and 'Filter Settings'. In the 'Enter New Filter' section, there is a text input field for 'Filter Name:'. Below that, there are two radio button options for 'Filter Object:'. The first option, 'Object ID:', is selected and has an adjacent text input field. The second option, 'Object ID Selection List:', is unselected and has a dropdown menu showing a list of MIB nodes: 'system', 'interfaces', 'ip', and 'icmp'. In the 'Filter Settings' section, there are two radio button options for 'Object Filter:'. The first option, 'Include Object', is selected, and the second option, 'Exclude Object', is unselected. At the bottom of the form, there are two buttons: 'Apply' (in a blue box) and 'Close' (in a grey box).

- Filter Name—Enter a name between 0-30 characters.
 - Filter Object—Select the node in the MIB tree that is included or excluded in the selected SNMP filter. The options to select the object are as follows:
 - Selection List—Enables you to navigate the MIB tree. Press the Up arrow to go to the level of the selected node's parent and siblings; press the Down arrow to descend to the level of the selected node's children. Click nodes in the view to pass from one node to its sibling. Use the scrollbar to bring siblings in view.
 - If Object ID is used, the object identifier is included in the view if the *Include in filter option* is selected.
4. Include or exclude in Object Filter. If this is selected, the selected MIBs are included in the filter, otherwise they are excluded.
 5. Click **Apply**. The SNMP views are defined and the running configuration is updated.

V1/V2 Notification Recipients

The screenshot shows the configuration interface for a Linksys LGS528 28-Port Gigabit Managed Switch. The top navigation bar includes 'System Status', 'Quick Start', 'Configuration', 'Maintenance', and 'Support'. The left sidebar lists various system management options, with 'V1/V2 Notification Recipients' selected under the 'SNMP' section. The main content area displays the 'V1/V2 Notification Recipient Table' with columns for Recipient IP Address, UDP Port, Notification Version, Notification Type, Community, and Filter Name. The table currently shows 0 results found. Below the table are buttons for 'Add', 'Edit', and 'Delete'. The footer contains the copyright notice: © 2014 Belkin International, Inc. and/or its subsidiaries and affiliates, including Linksys, LLC. All rights reserved.

Trap messages are generated to report system events, as defined in RFC 1215. The system can generate traps defined in the MIB that it supports.

Trap receivers (aka Notification Recipients) are network nodes where the trap messages are sent by the device. A list of notification recipients are defined as the targets of trap messages.

A trap receiver entry contains the IP address of the node and the SNMP credentials corresponding to the version that is included in the trap message. When an event arises that requires a trap message to be sent, it is sent to every node listed in the Notification Recipient Table.

The Notification Recipients SNMPv1/v2 page and the Notification Recipients SNMPv3 page enable configuring the destination to which SNMP notifications are sent, and the types of SNMP notifications that are sent to each destination (traps or informs). The Add/Edit pop-ups enable configuring the attributes of the notifications.

An SNMP notification is a message sent from the device to the SNMP management station indicating that a certain event has occurred, such as a link up/ down.

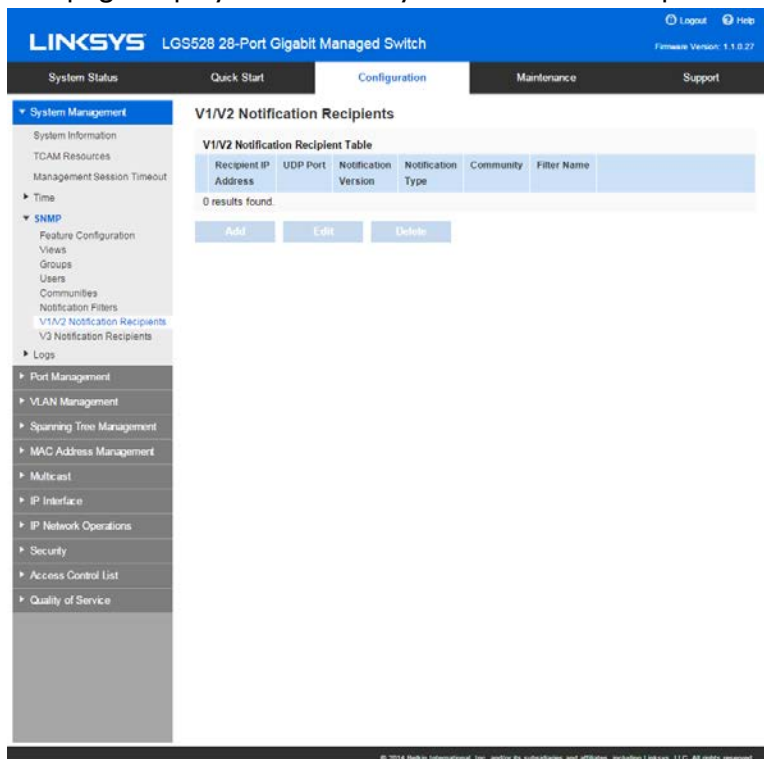
It is also possible to filter certain notifications. This can be done by creating a filter in the *Notification Filter* page and attaching it to an SNMP notification recipient. The notification filter enables filtering the type of SNMP notifications that are sent to the management station based on the OID of the notification that is about to be sent.

Defining SNMP Notification Recipients

To define a recipient in SNMPv1/v2:

1. Click *Configuration > System Management > SNMP > V1/V2 Notification Recipients*.

This page displays the currently-defined SNMP recipients.



2. Enter the parameters.

- Recipient—Select whether to specify the remote log server by IP address or server name.

- IP Version—Select either IPv4 or IPv6.
 - IPv6 Address Type—Select either Link Local or Global.
 - 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 whether it is received through a VLAN or ISATAP.
 - Recipient IP Address—Enter the IP address of where the traps are sent.
 - Recipient IP Name—Enter the server name of where the traps are sent.
 - UDP Port—Enter the UDP port used for notifications on the recipient device.
 - Notification Type—Select whether to send Traps or Informs. If both are required, two recipients must be created.
 - Notification Version—Select the trap SNMP version 1 or 2.
 - Community—Select from the drop-down the community string of the trap manager. Community String names are generated from those listed in the *Community* page.
 - Notification Filter—Select to enable filtering the type of SNMP notifications sent to the management station. The filters are created in the *Notification Filter* page.
 - Filter Name—Select the SNMP filter that defines the information contained in traps (defined in the *Notification Filter* page).
3. Click **Apply**. The SNMP Notification Recipient settings are written to the Running Configuration file.

V3 Notification Recipients

LINKSYS LGS528 28-Port Gigabit Managed Switch Firmware Version: 1.1.0.27

System Status Quick Start Configuration Maintenance Support

System Management

- System Information
- TCAM Resources
- Management Session Timeout
- Time
- SNMP
 - Feature Configuration
 - Views
 - Groups
 - Users
 - Communities
 - Notification Filters
 - V1/V2 Notification Recipients
 - V3 Notification Recipients
- Logs
- Port Management
- VLAN Management
- Spanning Tree Management
- MAC Address Management
- Multicast
- IP Interface
- IP Network Operations
- Security
- Access Control List
- Quality of Service

V3 Notification Recipients

V3 Notification Recipient Table

| Recipient IP Address | UDP Port | Notification Type | User Name | Security Level | Filter Name | |
|----------------------|----------|-------------------|-----------|----------------|-------------|--|
| 0 results found. | | | | | | |

Add Edit Delete

© 2014 Belkin International, Inc. and/or its subsidiaries and affiliates, including Linksys, LLC. All rights reserved.

To define a recipient in SNMPv3:

1. Click *SNMP > V3 Notification Recipients SNMP*.

This page displays recipients for SNMPv3.

2. Enter the fields:

- Recipient—Select whether to specify the remote log server by IP address or server name
- IP Version—Select either IPv4 or IPv6.

- IPv6 Address Type—Select the IPv6 address type (if IPv6 is used). The options:
 - 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.
- Link Local Interface—Select the link local interface (if *IPv6 Address Type Link Local* is selected) from the drop-down list.
- Recipient IP Address/Name—Enter the IP address or server name of where the traps are sent.
- UDP Port—Enter the UDP port used to for notifications on the recipient device.
- Notification Type—Select whether to send traps or informs. If both are required, two recipients must be created.
- User Name—Select from the drop-down list the user to whom SNMP notifications are sent. In order to receive notifications, this user must be defined on the SNMP User page, and its engine ID must be remote.
- Security Level—Select how much authentication is applied to the packet.

Note—*The Security Level here depends on which User Name was selected. If this User Name was configured as No Authentication, the Security Level is No Authentication only. However, if this User Name has assigned Authentication and Privacy on the User page, the security level on this screen can be either No Authentication, or Authentication Only, or Authentication and Privacy.*

The options are:

- No Authentication—Indicates the packet is neither authenticated nor encrypted.
 - Authentication—Indicates the packet is authenticated but not encrypted.
 - Privacy—Indicates the packet is both authenticated and encrypted.
- Notification Filter—Select to enable filtering the type of SNMP notifications sent to the management station. The filters are created in the Notification Filter page.

- Filter Name—Select the SNMP filter that defines the information contained in traps (defined in the *Notification Filter* page).
3. Click **Apply**. The SNMP Notification Recipient settings are written to the Running Configuration file.

Logs

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

- o [Overview](#)
- o [Log Management](#)
- o [Remote Log Servers](#)
- o [RAM Log](#)
- o [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.

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, which 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:

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

LINKSYS LGS528 28-Port Gigabit Managed Switch Logout Help
Firmware Version: 1.1.0.27

System Status Quick Start Configuration Maintenance Support

System Management

- System Information
- TCAM Resources
- Management Session Timeout
- Time
- SNMP
- Logs**
 - Log Management
 - Remote Log Servers
 - RAM Log
 - Flash Memory Log
- Port Management
- VLAN Management
- Spanning Tree Management
- MAC Address Management
- Multicast
- IP Interface
- IP Network Operations
- Security
- Access Control List
- Quality of Service

Log Management

SYSTEM LOG

Logging: Enable

Originator Identifier: None Hostname
 IPv4 Address IPv6 Address
 User Defined

LOG SETTINGS

| Severity Level | RAM Logging | Flash Memory Logging |
|----------------|-------------------------------------|-------------------------------------|
| Emergency: | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
| Alert: | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
| Critical: | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
| Error: | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
| Warning: | <input checked="" type="checkbox"/> | <input type="checkbox"/> |
| Notice: | <input checked="" type="checkbox"/> | <input type="checkbox"/> |
| Informational: | <input checked="" type="checkbox"/> | <input type="checkbox"/> |
| Debug: | <input type="checkbox"/> | <input type="checkbox"/> |

Apply Cancel

© 2014 Belkin International, Inc. and/or its subsidiaries and affiliates, including Linksys, LLC. All rights reserved.

2. Enter the parameters.

- System Log
 - Logging—Select to enable message logging.
 - Originator Identifier—Enables adding an origin identifier to SYSLOG messages. The options:
 - o None—Do not include the origin identifier in SYSLOG messages.
 - o Hostname—Include the system hostname in SYSLOG messages.
 - o IPv4 Address—Include the IPv4 address of the sending interface in SYSLOG messages.
 - o IPv6 Address—Include the IPv6 address of the sending interface in SYSLOG messages.
 - o 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 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.

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:

1. Click *Configuration > System Management > Logs > Remote Log Servers*.
The list of configured remote log servers is displayed.
2. Click **Add**.
3. Enter the parameters.

- Enter New Server
 - 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:
 - o Global—The IPv6 address is a global Unicast IPV6 type that is visible and reachable from other networks.
 - o 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.
 - o 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.
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.

Chapter 5 - 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](#)
- [Discovery - LLDP](#)

Ports

To configure ports:

1. Configure port by using the Ports page.
2. Enable/disable the Link Aggregation Control (LAG) protocol, and configure the potential member ports to the desired LAGs by using the LAGs 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.
8. Configure LLDP and LLDP-MED port configuration by using Discovery LLDP.

To configure port settings:

1. Click *Configuration > Port Management > Ports*.
2. Select Enable to support jumbo packets of up to 9 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.
3. To update the port settings, select the desired port, and click Edit.

4. Modify the following parameters:

The screenshot shows the 'Edit Port' configuration interface. At the top, it says 'Edit Port'. Below that, there's a section 'Select Your Port' with a dropdown menu showing 'GE1'. Underneath is the 'Port Settings' section. It includes: 'Operational Status: 1000M-Copper Up 1000M Full'; 'Administrative Mode: Up (selected) Down Suspended Port: Reactivate'; 'Protected Port: Enable'; 'Auto Negotiation: Enable (checked)'; 'Port Speed: 10M 100M 1000M (selected)'; 'Duplex Mode: Half Full (selected)'; 'Auto Advertisement: Max Capability (checked) 10 Full Duplex 10 Half Duplex 100 Full Duplex 100 Half Duplex 1000 Full Duplex'; 'Back Pressure: Enable'; 'Flow Control: Enable Disable (selected) Auto Negotiation'; 'MDI/MDIX: MDIX MDI Auto (selected)'; and a 'Description:' text box. At the bottom, there are 'Apply' and 'Close' buttons.

- Port—Select the port number.
- 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:

- 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.

Possible options:

- 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.

Options:

- 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.

Options:

- 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.
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](#)

[Load Balancing](#)

[LAG Management](#)

[Default Settings and Configuration](#)

[Static and Dynamic LAG Workflow](#)

[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 eight 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:

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:

1. Enable LACP on the LAG. Assign up to 16 candidate 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:

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

LINKSYS LGS528 28-Port Gigabit Managed Switch Logout Help
Firmware Version: 1.1.0.27

System Status Quick Start **Configuration** Maintenance Support

System Management
 System Information
 TCAM Resources
 Management Session Timeout
 Time
 SNMP
 Logs

Port Management
 Ports
 Link Aggregation
 LAGs
 Green Ethernet
 Discovery - LLDP

VLAN Management
 Spanning Tree Management
 MAC Address Management
 Multicast
 IP Interface
 IP Network Operations
 Security
 Access Control List
 Quality of Service

LAGs

Load Balance Method: by MAC address by IP & MAC Addresses

Apply Cancel

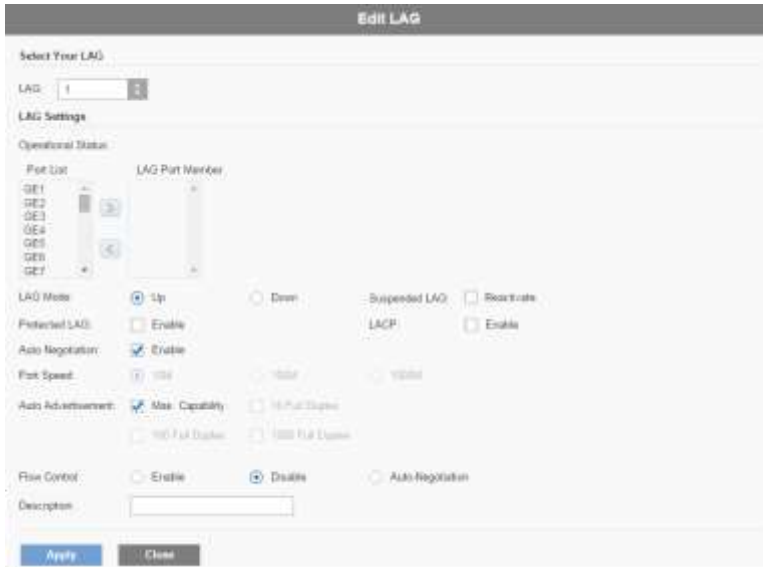
LAG Table

| | LAG | LAG Type | Operational Status | Active Port Members | Standby Port Members | LACP | Operational LAG Speed | Operational Flow Control | Protected LAG |
|----------------------------------|-------|----------|--------------------|---------------------|----------------------|------|-----------------------|--------------------------|---------------|
| <input checked="" type="radio"/> | LAG 1 | | | | | | | | Unprotected |
| <input type="radio"/> | LAG 2 | | | | | | | | Unprotected |
| <input type="radio"/> | LAG 3 | | | | | | | | Unprotected |
| <input type="radio"/> | LAG 4 | | | | | | | | Unprotected |
| <input type="radio"/> | LAG 5 | | | | | | | | Unprotected |
| <input type="radio"/> | LAG 6 | | | | | | | | Unprotected |
| <input type="radio"/> | LAG 7 | | | | | | | | Unprotected |
| <input type="radio"/> | LAG 8 | | | | | | | | Unprotected |

Edit

© 2014 Belkin International, Inc. and/or its subsidiaries and affiliates, including Linksys, LLC. All rights reserved.

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



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:
 - 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.

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.*

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 1 Gigabyte, 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.

1. Ensure that auto-negotiation is enabled on the port by opening the Ports page.
 - Select a port and open the Edit Ports page.
 - Select Auto Negotiation field to ensure that it is Enabled.

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.
3. Ensure that 802.3 EEE is enabled on a port by opening the Green Ethernet page.
 - Select a port, open the Edit Ports page.
 - 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:

1. Click *Configuration > Port Management > Green Ethernet*.
2. Choose whether to enable the following features:

Green Ethernet

SETTINGS

| | |
|--|--|
| Energy Detect Mode: | <input type="checkbox"/> Enable |
| Short Reach: | <input type="checkbox"/> Enable |
| Port LEDs: | <input checked="" type="checkbox"/> Enable |
| 802.3 Energy Efficient Ethernet (EEE): | <input checked="" type="checkbox"/> Enable |

Apply

Cancel

- Energy Detect Mode—Select to globally enable.
- 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.

- Port LEDs—Select to disable port LEDs. When disabled, ports do not display link status, activity, etc.
 - 802.3 Energy Efficient Ethernet (EEE)—Select to globally enable EEE.
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:

The screenshot shows the configuration page for the Green Ethernet feature on a Linksys LGS528 28-Port Gigabit Managed Switch. The page is divided into several sections:

- System Management:** System Status, Quick Start, Configuration, Maintenance, Support.
- Green Ethernet SETTINGS:**
 - Energy Detect Mode: Enable
 - Short Reach: Enable
 - Port LEDs: Enable
 - 802.3 Energy Efficient Ethernet (EEE): Enable
- POWER EFFICIENCY:**
 - Power Savings: 70 %
 - Cumulative Energy Saved: 0 Watt Hour
- Green Ethernet Table:** A table with columns for Port, Short Reach Mode, Short Reach Status, Short Reach Reason, Cable Length, EEE Mode, EEE Status, LLDP Mode, LLDP Status, Remote EEE Mode, Energy Detect Mode, Energy Detect Status, and Energy Detect Reason. The table lists 28 ports (GE1 to GE28) with their respective settings.

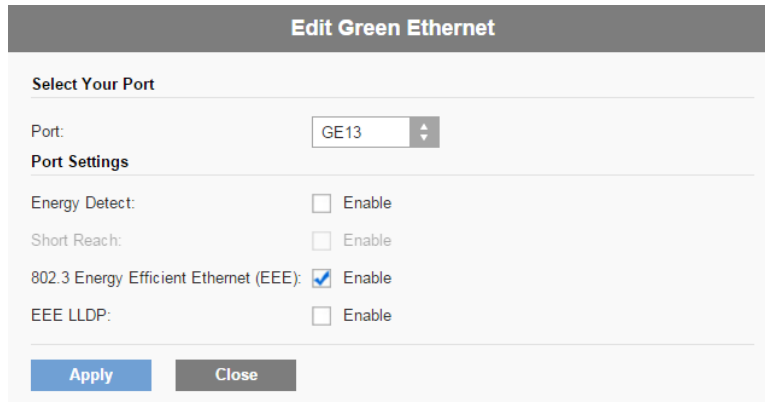
- Port—The port number.
- Short-Reach Mode—Whether Short-Reach Mode is enabled.
- Short Reach Status—Whether Short-Reach Mode is operational. This is a function of whether it has been enabled (Administrative Status), whether it has been enabled on the local port, and whether it is operational on the 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 of the port regarding to EEE mode (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.

4. Select a Port and click **Edit**.



5. Select to enable or disable the various features.
6. 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:

- 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:

- 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 device does 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

Feature Configuration

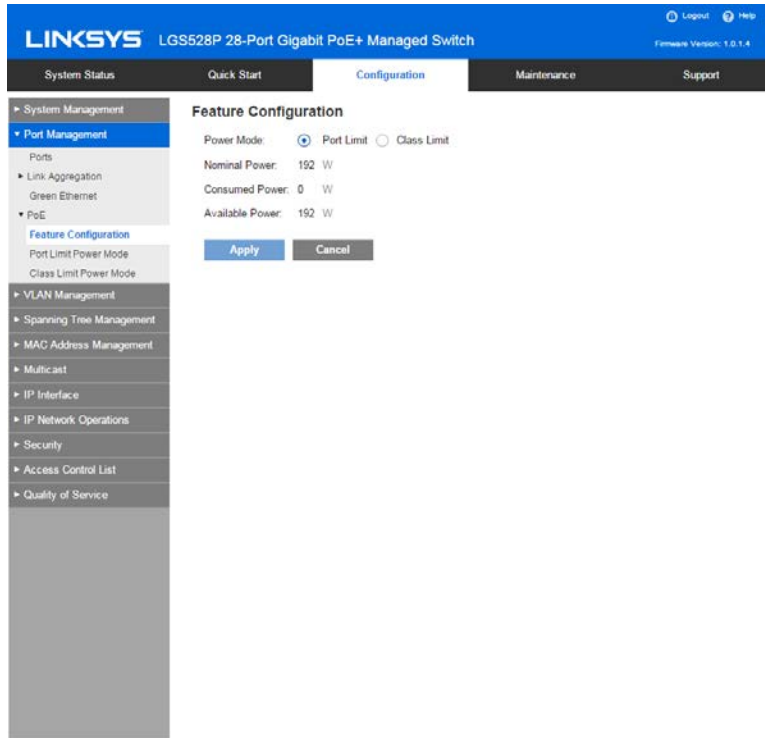
The Feature Configuration page enables selecting either the Port Limit or Class Limit PoE mode and specifying the PoE traps to be generated.

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:

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



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.
3. Click **Apply** to save the PoE properties.

Port Limit Power Mode

To configure port limit power mode:

1. Click *Configuration > Port Management > PoE > Port Limit Power Mode*.

LINKSYS LGS528P 28-Port Gigabit PoE+ Managed Switch Legend Help
Firmware Version: 1.0.1.4

System Status Quick Start **Configuration** Maintenance Support

System Management
Port Management
 Ports
 Link Aggregation
 Green Ethernet
 PoE
 Feature Configuration
Port Limit Power Mode
 Class Limit Power Mode
 VLAN Management
 Spanning Tree Management
 MAC Address Management
 Multicast
 IP Interface
 IP Network Operations
 Security
 Access Control List
 Quality of Service

Port Limit Power Mode

| Port | PoE Status | Power Priority Level | Power Allocation Limit (mW) | Max Power Allocation (mW) | Power Consumption (mW) | Class | Operational Status |
|----------------------------|------------|----------------------|-----------------------------|---------------------------|------------------------|-------|--------------------|
| <input type="radio"/> GE1 | Enabled | Low | 30000 | 31500 | 0 | 4 | Searching |
| <input type="radio"/> GE2 | Enabled | Low | 30000 | 31500 | 0 | 4 | Searching |
| <input type="radio"/> GE3 | Enabled | Low | 30000 | 31500 | 0 | 4 | Searching |
| <input type="radio"/> GE4 | Enabled | Low | 30000 | 31500 | 0 | 4 | Searching |
| <input type="radio"/> GE5 | Enabled | Low | 30000 | 31500 | 0 | 4 | Searching |
| <input type="radio"/> GE6 | Enabled | Low | 30000 | 31500 | 0 | 4 | Searching |
| <input type="radio"/> GE7 | Enabled | Low | 30000 | 31500 | 0 | 4 | Searching |
| <input type="radio"/> GE8 | Enabled | Low | 30000 | 31500 | 0 | 4 | Searching |
| <input type="radio"/> GE9 | Enabled | Low | 30000 | 31500 | 0 | 4 | Searching |
| <input type="radio"/> GE10 | Enabled | Low | 30000 | 31500 | 0 | 4 | Searching |
| <input type="radio"/> GE11 | Enabled | Low | 30000 | 31500 | 0 | 4 | Searching |
| <input type="radio"/> GE12 | Enabled | Low | 30000 | 31500 | 0 | 4 | Searching |
| <input type="radio"/> GE13 | Enabled | Low | 30000 | 31500 | 0 | 4 | Searching |
| <input type="radio"/> GE14 | Enabled | Low | 30000 | 31500 | 0 | 4 | Searching |
| <input type="radio"/> GE15 | Enabled | Low | 30000 | 31500 | 0 | 4 | Searching |
| <input type="radio"/> GE16 | Enabled | Low | 30000 | 31500 | 0 | 4 | Searching |
| <input type="radio"/> GE17 | Enabled | Low | 30000 | 31500 | 0 | 4 | Searching |
| <input type="radio"/> GE18 | Enabled | Low | 30000 | 31500 | 0 | 4 | Searching |
| <input type="radio"/> GE19 | Enabled | Low | 30000 | 31500 | 0 | 4 | Searching |
| <input type="radio"/> GE20 | Enabled | Low | 30000 | 31500 | 0 | 4 | Searching |
| <input type="radio"/> GE21 | Enabled | Low | 30000 | 31500 | 0 | 4 | Searching |
| <input type="radio"/> GE22 | Enabled | Low | 30000 | 31500 | 0 | 4 | Searching |
| <input type="radio"/> GE23 | Enabled | Low | 30000 | 31500 | 0 | 4 | Searching |
| <input type="radio"/> GE24 | Enabled | Low | 30000 | 31500 | 0 | 4 | Searching |

[Edit](#)

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.
- Select a port and click **Edit**. Enter the fields as described above.

Edit Port Limit

Select Your Port

Port: GE1

Port Settings

PoE Status: Enable

Power Priority Level: Critical High Low

Power Allocation Limit: 30000 mW (0-30000)

Port Status

Class: 4

Max Power Allocation: 31500 mW

Power Consumption: 0 mW

Apply Close

- 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:

- Click *Configuration > Port Management > PoE > Class Limit Power Mode*.

LINKSYS LGS528P 28-Port Gigabit PoE+ Managed Switch Logout Help
 Firmware Version: 1.0.1.4

System Status Quick Start **Configuration** Maintenance Support

- System Management
- Port Management**
 - Ports
 - Link Aggregation
 - Green Ethernet
 - PoE
 - Feature Configuration
 - Port Limit Power Mode
 - Class Limit Power Mode**
- VLAN Management
- Spanning Tree Management
- MAC Address Management
- Multicast
- IP Interface
- IP Network Operations
- Security
- Access Control List
- Quality of Service

Class Limit Power Mode

| Class Limit Table | | | | | | | |
|----------------------------|------------|----------------------|-------|---------------------------|------------------------|--------------------|--|
| Port | PoE Status | Power Priority Level | Class | Max Power Allocation (mW) | Power Consumption (mW) | Operational Status | |
| <input type="radio"/> GE1 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE2 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE3 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE4 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE5 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE6 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE7 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE8 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE9 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE10 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE11 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE12 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE13 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE14 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE15 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE16 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE17 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE18 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE19 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE20 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE21 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE22 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE23 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE24 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE25 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE26 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE27 | Enabled | Low | 4 | 30000 | 0 | Searching | |
| <input type="radio"/> GE28 | Enabled | Low | 4 | 30000 | 0 | Searching | |

[Edit](#)

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. The switch hardware may actually supply 5-10% more power than Max Power Allocation to accommodate the power loss over the wire.
- 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.

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

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

Discovery - LLDP

This section provides information for configuring Discovery. It covers the following topics:

- [Overview](#)
- [Feature Configuration](#)
- [LLDP MED Ports](#)
- [LLDP Local Information](#)
- [LLDP Neighbor Information](#)
- [LLDP MED Network Policy](#)

Overview

Link Layer Discovery Protocol (LLDP) is a link layer protocol for directly-connected LLDP-capable neighbors to advertise themselves and their capabilities. LLDP enables network managers to troubleshoot and enhance network management in multi-vendor environments. LLDP standardizes methods for network devices to advertise themselves to other systems, and to store discovered information.

By default, the device sends an LLDP advertisement periodically to all its interfaces and processes incoming LLDP packets as required by the protocols. In LLDP, advertisements are encoded as TLV (Type, Length, Value) in the packet.

The information learned is stored in the data in a Management Information Base (MIB). The network management system models the topology of the network by querying these MIB databases.

By default, the device terminates and processes all incoming LLDP packets as required by the protocol.

The LLDP protocol has an extension called LLDP Media Endpoint Discovery (LLDP-MED) that provides and accepts information from media endpoint devices such as VoIP phones and video phones. For further information about LLDP-MED, see LLDP MED Network Policy.

The following LLDP configuration notes apply:

- LLDP can be enabled or disabled globally or per port. The LLDP capability of a port is relevant only if LLDP is globally enabled.
- If LLDP is globally enabled, the device filters out incoming LLDP packets from ports that are LLDP-disabled.
- If LLDP is globally disabled, the device can be configured to discard, VLAN-aware flooding, or VLAN-unaware flooding of all incoming LLDP packets. VLAN-aware flooding floods an incoming LLDP packet to the VLAN where the packet is received excluding the ingress port. VLAN-unaware flooding floods an incoming LLDP packet to all the ports excluding the ingress port. The default is to discard LLDP packets when LLDP is globally disabled. You can configure the discard/flooding of incoming LLDP packets from the LLDP Feature Configuration page.

LLDP end devices, such as IP phones, learn the voice VLAN configuration from LLDP advertisements. By default, the device is enabled to send out LLDP advertisement based on the voice VLAN configured at the device. Refer to the Voice VLAN ([p. 118](#)) for details.

Note--*LLDP does not distinguish if a port is in a LAG. If there are multiple ports in a LAG, LLDP transmit packets on each port without taking into account the fact that the ports are in a LAG.*

The operation of LLDP is independent of the STP status of an interface.

If 802.1x port access control is enabled at an interface, the device transmits and receives LLDP packets to and from the interface only if the interface is authenticated and authorized.

If a port is the target of mirroring, then LLDP considers it down.

Note—LLDP are link layer protocols for directly-connected LLDP capable devices to advertise themselves and their capabilities. In deployments where the LLDP-capable devices are not directly connected and are separated with LLDP-incapable devices, the LLDP-capable devices may be able to receive the advertisement from other device(s) only if the LLDP-incapable devices flood the LLDP packets they receive. If the LLDP-incapable devices perform VLAN-aware flooding, then LLDP-capable devices can hear each other only if they are in the same VLAN. An LLDP-capable device may receive advertisements from more than one device if the LLDP-incapable devices flood the LLDP packets.

Workflows

Following are examples of actions that can be performed with the LLDP feature and in a suggested order. You can refer to the LLDP section for additional guidelines on LLDP configuration. LLDP configuration pages are accessible under the *Administration Configuration > Port Management > Discovery LLDP* menu.

1. Enter LLDP global parameters, such as LLDP Frames Handling using the LLDP Feature Configuration page.
2. Configure LLDP per port by using LLDP Feature Configuration page. On this page, interfaces can be configured to receive/transmit LLDP PDUs, send SNMP notifications, specify which TLVs to advertise, and advertise the device's management address.
3. Create LLDP MED network policies by using the LLDP MED Network Policy page.
4. Associate LLDP MED network policies and the optional LLDP-MED TLVs to the desired interfaces by using the LLDP MED Port Settings page.

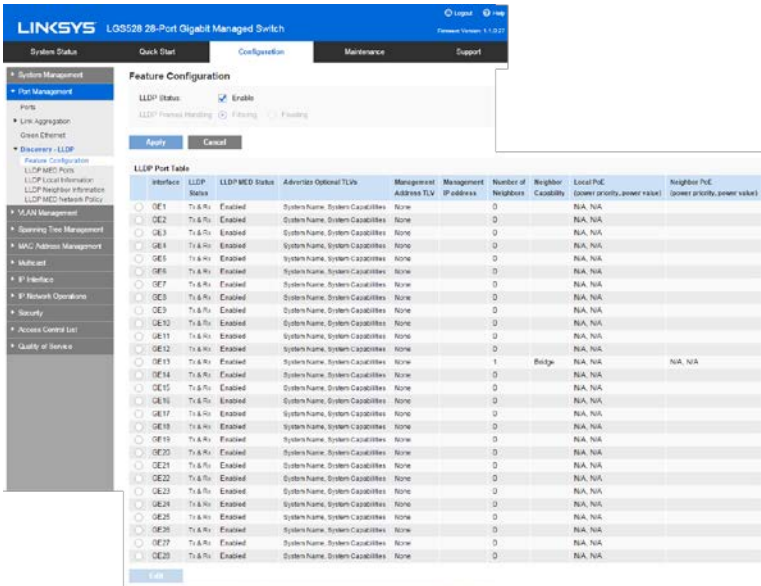
Feature Configuration

The Feature Configuration Page enables configuration of LLDP global parameters and entering the TLVs that are sent in the LLDP PDU.

The LLDP-MED TLVs to be advertised can be selected in the LLDP MED Port Settings page, and the management address TLV of the device may be configured to be advertised.

To configure the LLDP port settings:

1. Click *Configuration > Port Management > Discovery - LLDP > Feature Configuration*.



The following fields are displayed (only fields that do not appear in the Edit page are described):

- Interface—The port to edit.
- LLDP MED Status—Enabled or disabled.
- Number of neighbors—Number of neighbors discovered.
- Neighbor Capability—Displays the primary functions of the neighbor; for example: Bridge or Router.
- Local PoE—Local PoE information advertised.
 - power priority—Port power priority
 - power value—Port power value
- Neighbor PoE—PoE information advertised by the neighbor.
 - power priority—Port power priority
 - power value—Port power value

2. Enter the following fields.

- LLDP Status—Select to enable LLDP on the device (enabled by default).
- LLDP Frame Handling—If LLDP is not enabled, select the action to be taken if a packet that matches the selected criteria is received.
 - Filtering—Delete the packet.
 - Flooding—Forward the packet to all VLAN members.

3. Select a port and click **Edit**.

This page provides the following fields:

- Port—Select the port to edit.
- LLDP Status—Select the LLDP publishing option for the port. The values are the following:
 - Tx Only—Publishes but does not discover.
 - Rx Only—Discovers but does not publish.
 - Tx & Rx—Publishes and discovers.
 - Disable—Indicates that LLDP is disabled on the port.
- Management Address TLV—Select one of the following ways to advertise the IP management address of the device:
 - Auto Advertise—Specifies that the software automatically chooses a management address to advertise from all the IP addresses of the device. In case of multiple IP addresses, the software chooses the lowest IP address among the dynamic IP addresses. If there are no dynamic addresses, the software chooses the lowest IP address among the static IP addresses.
 - Manual Advertise—Select this option and the management IP address to be advertised.
 - None—Do not advertise the management IP address.
- Management IP Address—If Manual Advertise was selected, select the Management IP address from the addresses provided.
- Available Optional TLVs—Information to be published by the device

- Advertise Optional TLVs—Select the information to be published by the device by moving the TLV from the Available Optional TLVs list. The available TLVs contain the following information:
 - Port Description—Information about the port, including manufacturer, product name and hardware/software version.
 - System Name—System's assigned name (in alpha-numeric format). The value equals the sysName object.
 - System Description—Description of the network entity (in alpha-numeric format). This includes the system's name and versions of the hardware, operating system, and networking software supported by the device. The value equals the sysDescr object.
 - System Capabilities—Primary functions of the device, and whether or not these functions are enabled on the device. The capabilities are indicated by two octets. Bits 0 through 7 indicate Other, Repeater, Bridge, WLAN AP, Router, Telephone, DOCSIS cable device, and station respectively. Bits 8 through 15 are reserved.
 - 802.3 MAC-PHY—Duplex and bit rate capability and the current duplex and bit rate settings of the sending device. It also indicates whether the current settings are due to auto-negotiation or manual configuration.
 - 802.3 Link Aggregation—Whether the link (associated with the port on which the LLDP PDU is transmitted) can be aggregated. It also indicates whether the link is currently aggregated, and if so, provides the aggregated port identifier.
 - 802.3 Maximum Frame—Maximum frame size capability of the MAC/PHY
4. Enter the relevant information, and click **Apply**. The port settings are written to the Running Configuration file

LLDP MED Ports

The LLDP MED Ports page enables the selection of the LLDP-MED TLVs and/or the network policies to be included in the outgoing LLDP advertisement for the desired interfaces. Network Policies are configured using the LLDP MED Network Policy page.

To configure LLDP MED on each port:

1. Click *Configuration > Port Management > Discovery - LLDP > LLDP MED Ports*.

This page displays the following LLDP MED settings for all ports (only fields not described in the Edit page are listed):

- Location—Whether Location TLV is transmitted.
 - PoE—Whether POE-PSE TLV is transmitted.
 - Inventory—Whether Inventory TLV is transmitted.
2. The message at the top of the page indicates whether the generation of the LLDP MED Network Policy for the voice application is automatic or not (see LLDP Overview). Click on the link to change the mode.
 3. To associate additional LLDP MED TLV and/or one or more user-defined LLDP MED Network Policies to a port, select it, and click **Edit**.
 4. Enter the parameters:

Edit LLDP MED Port

Select Your Port

Port: GE13

Port Settings

LLDP MED Status: Enable

Available Optional TLVs: Location, Inventory

Advertise Optional TLVs: Network Policy

Available Network Policies:

Advertise Network Policies:

Location Coordinate: (16 pairs of hex digits)

Location Civic Address: (6-160 pairs of hex digits)

Location ECS ELIN: (10-25 pairs of hex digits)

Apply Close

- Port—Select the interface to configure.
- LLDP MED Status—Enable/disable LLDP MED on this port.
- Available Optional TLVs—Select the TLVs that can be published by the device by moving them from the Advertise Optional TLVs list.

- Available Network Policies—Select the LLDP MED policies to be published by LLDP by moving them from the Available Network Policies list. These were created in the LLDP MED Network Policy page. To include one or more user-defined network policies in the advertisement, you must also select Network Policy from the Available Optional TLVs.

Note—*The following fields must be entered in hexadecimal characters in the exact data format that is defined in the LLDP-MED standard (ANSI-TIA-1057_final_for_publication.pdf):*

- Location Coordinate—Enter the coordinate location to be published by LLDP.
 - Location Civic Address—Enter the civic address to be published by LLDP.
 - Location (ECS) ELIN—Enter the Emergency Call Service (ECS) ELIN location to be published by LLDP.
5. Click **Apply**. The LLDP MED port settings are written to the Running Configuration file.

LLDP Local Information

To view the LLDP local port status advertised on a port:

1. Click *Configuration > Port Management Discovery - LLDP > LLDP Local Information*.
2. Select the desired port from the Port list.

This page displays the following groups of fields (the actual fields displayed depend on the optional TLVs selected to be advertised):

The screenshot shows the configuration page for a Linksys LGS528 28-Port Gigabit Managed Switch. The page is titled "LLDP Local Information" and is divided into several sections:

- PORT:** A dropdown menu showing "GE1".
- LOCAL LLDP INFORMATION:**
 - Global:**

| | |
|--------------------------------|-------------------|
| Chassis ID Subtype: | MAC address |
| Chassis ID: | b4:75:0e:7c:5b:7a |
| System Name: | switch7c5b7a |
| System Description: | N/A |
| Supported System Capabilities: | Bridge, Router |
| Enabled System Capabilities: | Bridge, Router |
| Port ID Subtype: | Interface name |
| Port ID: | gi1 |
| Port Description: | gigabitethernet1 |
 - Management Address:**

| | |
|--------------------|-----|
| Address Subtype: | N/A |
| Address: | N/A |
| Interface Subtype: | N/A |
| Interface Number: | N/A |
- LOCAL LLDP MED INFORMATION:**
 - MED Information:**

| | |
|-------------------------|-----|
| Capabilities Supported: | N/A |
| Current Capabilities: | N/A |
| Device Class: | N/A |
| PoE Device Type: | N/A |
| PoE Power Source: | N/A |
| PoE Power Priority: | N/A |
| PoE Power Value: | N/A |
| Hardware Revision: | N/A |
| Firmware Revision: | N/A |
| Software Revision: | N/A |
| Serial Number: | N/A |
| Manufacturer Name: | N/A |
| Model Name: | N/A |
| Asset ID: | N/A |
 - Location Information:**

| | |
|--------------|-----|
| Civic: | N/A |
| Coordinates: | N/A |
| ECS ELIN: | N/A |
 - Network Policy:**

| Application Type | VLAN ID | VLAN Type | User Priority | DSCP |
|------------------|---------|-----------|---------------|------|
| 0 results found. | | | | |

© 2014 Belkin International, Inc. and/or its subsidiaries and affiliates, including Linksys, LLC. All rights reserved.

- Global
 - Chassis ID Subtype—Type of chassis ID. (For example, the MAC address.)
 - Chassis ID—Identifier of chassis. Where the chassis ID subtype is a MAC address, the MAC address of the device appears.
 - System Name—Name of device.
 - System Description—Description of the device (in alpha-numeric format).
 - Supported System Capabilities—Primary functions of the device, such as Bridge, WLAN AP, or Router.
 - Enabled System Capabilities—Primary enabled function(s) of the device.

- Port ID Subtype—Type of the port identifier that is shown.
- Port ID—Identifier of port.
- Port Description—Information about the port, including manufacturer, product name and hardware/software version.
- Management Address

Displays the table of addresses of the local LLDP agent. Other remote managers can use this address to obtain information related to the local device. The address consists of the following elements:

 - Address Subtype—Type of management IP address that is listed in the Management Address field; for example, IPv4.
 - Address—Returned address most appropriate for management use.
 - Interface Subtype—Numbering method used for defining the interface number.
 - Interface Number—Specific interface associated with this management address.
- MED Information
 - Capabilities Supported—MED capabilities supported on the port.
 - Current Capabilities—MED capabilities enabled on the port.
 - Device Class—LLDP-MED endpoint device class. The possible device classes are:
 - Endpoint Class 1—Generic endpoint class, offering basic LLDP services.
 - Endpoint Class 2—Media endpoint class, offering media streaming capabilities, as well as all Class 1 features.
 - Endpoint Class 3—Communications device class, offering all Class 1 and Class 2 features plus location, 911, Layer 2 device support, and device information management capabilities.
 - PoE Device Type—Port PoE type; for example, powered.
 - PoE Power Source—Port power source.
 - PoE Power Priority—Port power priority.
 - PoE Power Value—Port power value.
 - Hardware Revision—Hardware version.
 - Firmware Revision—Firmware version.
 - Software Revision—Software version.

- Serial Number—Device serial number.
- Manufacturer Name—Device manufacturer name.
- Model Name—Device model name.
- Asset ID—Asset ID.
- Location Information
 - Civic—Street address.
 - Coordinates—Map coordinates: latitude, longitude, and altitude.
 - ECS ELIN—Emergency Call Service (ECS) Emergency Location Identification Number (ELIN).
- Network Policy
 - Application Type—Network policy application type; for example, Voice.
 - VLAN ID—VLAN ID for which the network policy is defined.
 - VLAN Type—VLAN type for which the network policy is defined. The possible field values are the following:
 - Tagged—Indicates the network policy is defined for tagged VLANs.
 - Untagged—Indicates the network policy is defined for untagged VLANs.
 - User Priority—Network policy user priority.
 - DSCP—Network policy DSCP.

LLDP Neighbor Information

The LLDP Neighbor Information page contains information that was received from neighboring devices.

After timeout (based on the value received from the neighbor Time To Live TLV during which no LLDP PDU was received from a neighbor), the information is deleted.

To view the LLDP neighbors information:

1. Click *Configuration > Port Management > Discovery - LLDP > LLDP Neighbor Information*.

The screenshot shows the configuration page for a Linksys LGS528 28-Port Gigabit Managed Switch. The page is titled "LLDP Neighbor Information" and is part of the "Configuration" section. The left sidebar shows a navigation menu with "Port Management" selected, and "Discovery - LLDP" expanded to show "LLDP Neighbor Information".

| LLDP Neighbor Information | |
|--------------------------------|-------------------|
| PORT | |
| Port: | GE13 |
| NEIGHBOR LLDP INFORMATION | |
| Global | |
| MSAP Entry: | 7 |
| Chassis ID Subtype: | MAC Address |
| Chassis ID: | b4:75:0e:2d:98:9e |
| Port ID Subtype: | Interface Name |
| Port ID: | gi5 |
| Port Description: | |
| System Name: | switch2d989e |
| System Description: | |
| Supported System Capabilities: | Bridge |
| Enabled System Capabilities: | Bridge |

This page contains the following fields:

- Port—Number of the local port to which the neighbor is connected.
- Global
 - Local Port—Port number.
 - MSAP Entry—Device Media Service Access Point (MSAP) entry number.
 - Basic Details
 - Chassis ID Subtype—Type of chassis ID (for example, MAC address).
 - Chassis ID—Identifier of the 802 LAN neighboring device chassis.
 - Port ID Subtype—Type of the port identifier that is shown.
 - Port ID—Identifier of port.
 - Port Description—Information about the port, including manufacturer, product name and hardware/software version.
 - System Name—Name of system that is published.
 - System Description—Description of the network entity (in alpha-numeric format). This includes the system name and versions of the hardware, operating system, and networking software supported by the device. The value equals the sysDescr object.
 - Supported System Capabilities—Primary functions of the device. The capabilities are indicated by two octets. Bits 0 through 7 indicate Other, Repeater, Bridge, WLAN AP, Router, Telephone, DOCSIS cable device, and station, respectively. Bits 8 through 15 are reserved.
 - Enabled System Capabilities—Primary enabled function(s) of the device.
- Management Address

- Address Subtype—Managed address subtype; for example, MAC or IPv4.
- Address—Managed address.
- Interface Subtype—Port subtype.
- Interface Number—Port number.
- MED Information
 - Capabilities Supported—MED capabilities enabled on the port.
 - Current Capabilities—MED TLVs advertised by the port.
 - Device Class—LLDP-MED endpoint device class. The possible device classes are:
 - Endpoint Class 1—Indicates a generic endpoint class, offering basic LLDP services.
 - Endpoint Class 2—Indicates a media endpoint class, offering media streaming capabilities as well as all Class 1 features.
 - Endpoint Class 3—Indicates a communications device class, offering all Class 1 and Class 2 features plus location, 911, Layer 2 switch support and device information management capabilities.
 - PoE Device Type—Port PoE type, for example, powered.
 - PoE Power Source—Port’s power source.
 - PoE Power Priority—Port’s power priority.
 - PoE Power Value—Port’s power value.
 - Hardware Revision—Hardware version.
 - Firmware Revision—Firmware version.
 - Software Revision—Software version.
 - Serial Number—Device serial number.
 - Manufacturer Name—Device manufacturer name.
 - Model Name—Device model name.
 - Asset ID—Asset ID.
- Location Information

Enter the following data structures in hexadecimal as described in section 10.2.4 of the ANSI-TIA-1057 standard:

 - Civic—Civic or street address.
 - Coordinates—Location map coordinates—latitude, longitude, and altitude.

- ECS ELIN—Device's Emergency Call Service (ECS) Emergency Location Identification Number (ELIN).
- Unknown—Unknown location information.
- Network Policy
 - Application Type—Network policy application type, for example, Voice.
 - VLAN ID—VLAN ID for which the network policy is defined.
 - VLAN Type—VLAN type, Tagged or Untagged, for which the network policy is defined.
 - User Priority—Network policy user priority.
 - DSCP—Network policy DSCP.

LLDP MED Network Policy

LLDP Media Endpoint Discovery (LLDP-MED) is an extension of LLDP that provides the following additional capabilities to support media endpoint devices:

- Enables the advertisement and discovery of network policies for real-time applications such as voice and/or video.
- Enables discovery of the device location to allow creation of location databases and, in the case of Voice over Internet Protocol (VoIP), Emergency Call Service (E-911) by using IP Phone location information.

LLDP MED sends Troubleshooting information alerts to network managers upon:

- Port speed and duplex mode conflicts
- QoS policy misconfigurations

Setting LLDP MED Network Policy

An LLDP-MED network policy is a related set of configuration settings for a specific real-time application such as voice, or video. A network policy, if configured, can be included in the outgoing LLDP packets to the attached LLDP media endpoint device. The media endpoint device must send its traffic as specified in the network policy it receives. For example, a policy can be created for VoIP traffic that instructs VoIP phone to:

- Send voice traffic on VLAN 10 as tagged packet and with 802.1p priority 5.
- Send voice traffic with DSCP 46.

Network policies are associated with ports by using the LLDP MED Port Settings page. An administrator can manually configure one or more network policies and the interfaces where the policies are to be sent. It is the administrator's responsibility to manually create the VLANs and their port memberships according to the network policies and their associated interfaces.

In addition, an administrator can instruct the device to automatically generate and advertise a network policy for voice application based on the voice VLAN maintained by the device. Refer to the Auto Voice VLAN section for details on how the device maintains its voice VLAN.

To define an LLDP MED network policy:

1. Click *Configuration > Port Management > Discovery - LLDP > LLDP MED Network*.
This page contains previously-created network policies.
2. When Network Policy for Voice Application is enabled, the device automatically generates and advertises a network policy with the current voice VLAN configuration. Go to *Voice VLAN > Feature Configuration* page to configure the voice VLAN.
3. Click **Apply** to add this setting to the Running Configuration file.
4. To define a new policy, click **Add**.

Add LLDP MED Network Policy

New Network Policy

Network Policy Number: 1

Network Policy Settings

Application: Voice

VLAN ID: (1-4094)

VLAN Tag: Tagged Untagged

Layer 2 Priority: 0 DSCP: 0

Apply **Close**

5. Enter the values:
 - Network Policy Number—Select the number of the policy to be created.
 - Application—Select the type of application (type of traffic) for which the network policy is being defined.
 - VLAN ID—Enter the VLAN ID to which the traffic must be sent.
 - VLAN Tag—Select whether the traffic is Tagged or Untagged.
 - Layer 2 Priority—Select the traffic priority applied to traffic defined by this network policy. This is the CoS value.
 - DSCP Value—Select the DSCP value to associate with application data sent by neighbors. This informs them how they must mark the application traffic they send to the device.
6. Click **Apply**. The network policy is defined.

Note—You must manually configure the interfaces to include the desired manually-defined network policies for the outgoing LLDP packets using the LLDP MED Port Settings.

Chapter 6 - VLAN Management

This section covers the following topics:

[Overview](#)

[VLANs](#)

[Interfaces](#)

[VLAN Memberships](#)

[GVRP](#)

[VLAN Groups](#)

[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.

VLAN Description

Each VLAN is configured with a unique VLAN ID (VID) 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 only one untagged VLAN but can be a member of multiple tagged VLANs.

A port in VLAN Access mode can be part of only one VLAN. If it is in General or 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.

The frame is discarded at the ingress port if Ingress Filtering is enabled and the ingress port is not a member of the VLAN to which the packet belongs. A frame is regarded as priority-tagged only if the VID in its VLAN tag is 0.

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:

- 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. For more information about GVRP, refer to the GVRP Settings section.

Some VLANs can have additional roles, including:

- Voice VLAN: For more information refer to the Voice VLAN section.

- Guest VLAN: Set in the Edit VLAN Authentication page.
- Default VLAN: For more information refer to the Configuring Default VLAN Settings section.
- Management VLAN (in Layer 2-system-mode systems): For more information refer to the Layer 2 IP Addressing section.

QinQ

QinQ allows packets between sites of a customer network to be forwarded over a provider network. The device is a provider bridge that supports port-based c- tagged service interface where the customer network/site connects to.

With QinQ, the device adds a service VLAN tag known as Service Tag (S-tag) when forwarding customer traffic over the network. The S-tag is used to segregate traffic between various customers, while preserving the customer VLAN tag, which are known as C-tags.

Customer traffic is encapsulated with an S-tag with TPID 0x8100, regardless of whether it was originally c-tagged or untagged. The S-tag enables this traffic to be treated as an aggregate within a provider bridge network, where the bridging is based on the S-tag VID (S-VID) only.

The S-Tag is preserved while traffic is forwarded through the network service provider's infrastructure, and is later removed by an egress device.

An additional benefit of QinQ is that there is no need to configure customers' edge devices.

QinQ is enabled in the *VLAN Management > Interface Settings* page.

VLANs

This section describes the GUI pages used to configure various types of VLANs. This section describes the following processes:

- [VLAN Configuration Workflow](#)
- [Default VLAN Settings](#)
- [VLANs - Creating VLANs](#)
- [Interface Settings](#)
- [VLAN Membership](#)
- [GVRP](#)

VLAN Configuration Workflow

To configure VLANs:

1. If required, change the default VLAN as described in the Default VLAN Settings section.
2. Create the required VLANs as described in the VLANs - Creating VLANs section.
3. Set the desired VLAN-related configuration for ports as described in the Interface Settings section.
4. Assign interfaces to VLANs as described in the Port to VLAN section or the VLAN Memberships section.
5. View the current VLAN port membership for all the interfaces as described in the VLAN Memberships section.
6. If required, configure VLAN groups as described in the MAC-based Groups and Protocol-based VLANs sections.

Default VLAN Settings

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.
- RADIUS servers cannot assign the default VLAN to 802.1x supplicants by using Dynamic VLAN Assignment.

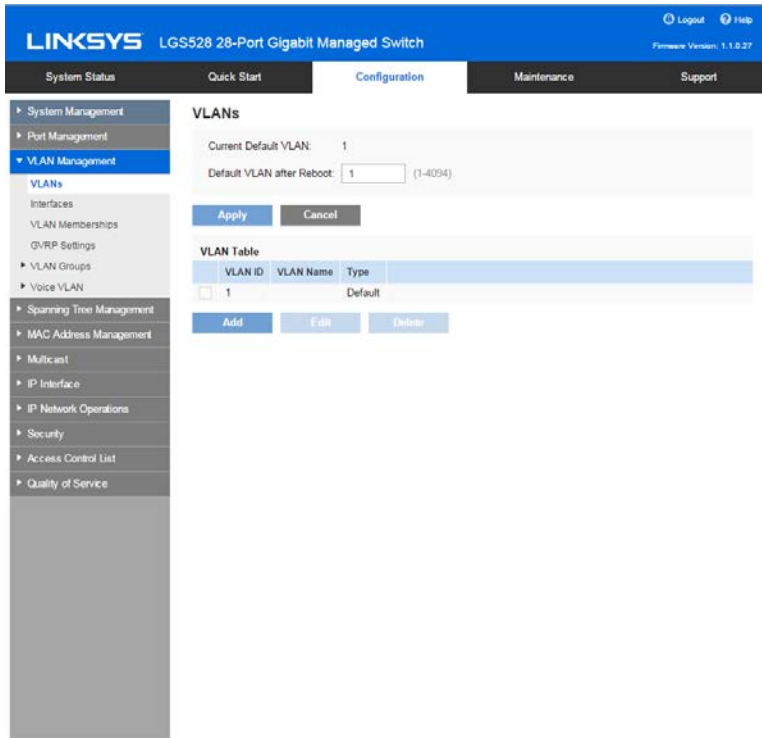
When the VID of the default VLAN is changed, the device performs the following on all the ports in the VLAN, after saving the configuration and rebooting the device:

- Removes VLAN membership of the ports from the original default VLAN (takes effect 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.

To change the default VLAN:

1. Click *Configuration > VLAN Management > VLANs*.



2. Enter the value for the following field:
 - Current Default VLAN ID—Displays the current default VLAN ID.
 - Default VLAN ID After Reboot—Enter a new VLAN ID to replace the default VLAN ID after reboot.
3. Click **Apply** to save the Running Configuration to the Startup Configuration.

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

VLANs - Creating VLANs

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 Managed device supports up to 41K VLANs, including the default VLAN.

Each VLAN must be configured with a unique VID with a value from 1 to 4094. The device reserves VID 4095 as the Discard VLAN and VID 4094 for 802.1x. All packets classified to the Discard VLAN are discarded at ingress, and are not forwarded to a port.

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

To change add a VLAN:

1. Click *Configuration > VLAN Management > VLANs*.
2. Click **Add** to add one or more new VLANs.

The screenshot shows a web interface for adding a VLAN. The title bar says "Add VLAN". Below it, there's a section "Enter New VLAN". It has two radio buttons: "Single VLAN" (which is selected) and "Range of VLANs". Underneath, there are three input fields: "VLAN ID:" with a range of "(2-4094)", "VLAN Name:", and "VLAN ID Range:" with a range of "(2-4094)". At the bottom, there are two buttons: "Apply" and "Close".

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

3. Enter the following fields for the new VLANs:
 - 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 .
4. Click **Apply** to create the VLAN(s).

Interfaces

LINKSYS LGS528 28-Port Gigabit Managed Switch Firmware Version: 1.1.0.27

System Status Quick Start Configuration Maintenance Support

System Management

- System Information
- TCAM Resources
- Management Session Timeout

Time

SNMP

Logs

Port Management

VLAN Management

- VLANs
- Interfaces**
- VLAN Memberships
- GVRP Settings
- VLAN Groups
- Voice VLAN

Spanning Tree Management

MAC Address Management

Multicast

IP Interface

IP Network Operations

Security

Access Control List

Quality of Service

Interfaces

Interface Table

Interface Type: Port Search

| Interface | Interface VLAN Mode | PVID | Acceptable Frame Type | Ingress Filtering | Administrative VLAN Memberships | Operational VLAN Memberships |
|----------------------------|---------------------|------|-----------------------|-------------------|---------------------------------|------------------------------|
| <input type="radio"/> GE1 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE2 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE3 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE4 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE5 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE6 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE7 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE8 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE9 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE10 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE11 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE12 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE13 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE14 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE15 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE16 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE17 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE18 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE19 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE20 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE21 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE22 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE23 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE24 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE25 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE26 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE27 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |
| <input type="radio"/> GE28 | Trunk | 1 | Admit All | Enabled | 1UP | 1UP |

Edit Join VLAN View

The Interface Settings page displays and enables configuration of VLAN-related parameters for all interfaces.

To configure the VLAN settings:

1. Click *Configuration*>*VLAN Management* > *Interfaces Settings*.
2. Select an interface type (Port or LAG), and click **Search**.
3. To configure a Port or LAG, select it and click **Edit**.

Note—To add a port or LAG to a VLAN, click *Join VLAN*.

4. Enter the values for the following fields:

The screenshot shows the 'Edit Interface' configuration page. It is divided into two main sections: 'Select Your Interface' and 'Interface Settings'. In the 'Select Your Interface' section, the 'Port' radio button is selected, and a dropdown menu shows 'GE13'. In the 'Interface Settings' section, the 'General Port' radio button is selected under 'Interface VLAN Mode'. The 'PVID' field contains the value '1'. Under 'Acceptable Frame Type', the 'Admit All' radio button is selected. The 'Ingress Filtering' checkbox is checked and labeled 'Enable'. At the bottom of the form are 'Apply' and 'Close' buttons.

- Interface—Select a Port/LAG.
- Interface VLAN Mode—Select the interface mode for the VLAN. The options are:
 - 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.
 - General Port—The interface can support all functions as defined in the IEEE 802.1q specification. The interface can be a tagged or untagged member of one or more VLANs.
- PVID—Enter the Port VLAN ID (PVID) of the VLAN to which incoming untagged and priority tagged frames are classified. The possible values are 1 to 4094.
- Acceptable Frame Type—Select the type of frame that the interface can receive. Frames that are not of the configured frame type are discarded at ingress. These frame types are only available in General mode.

Possible values are:

 - Admit All—The interface accepts all types of frames: untagged frames, tagged frames, and priority tagged frames.
 - Admit Tagged Only—The interface accepts only tagged frames.
 - Admit Untagged Only—The interface accepts only untagged and priority frames.
- Ingress Filtering—(Available only in General mode) Select to enable ingress filtering. When an interface is ingress filtering enabled, the interface discards all incoming frames that are classified as VLANs of which the interface is not a member. Ingress filtering can be disabled or enabled on general ports. It is always enabled on access ports and trunk ports.

5. Click **Apply**. The parameters are written to the Running Configuration file.

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.

To add a port to a VLAN:

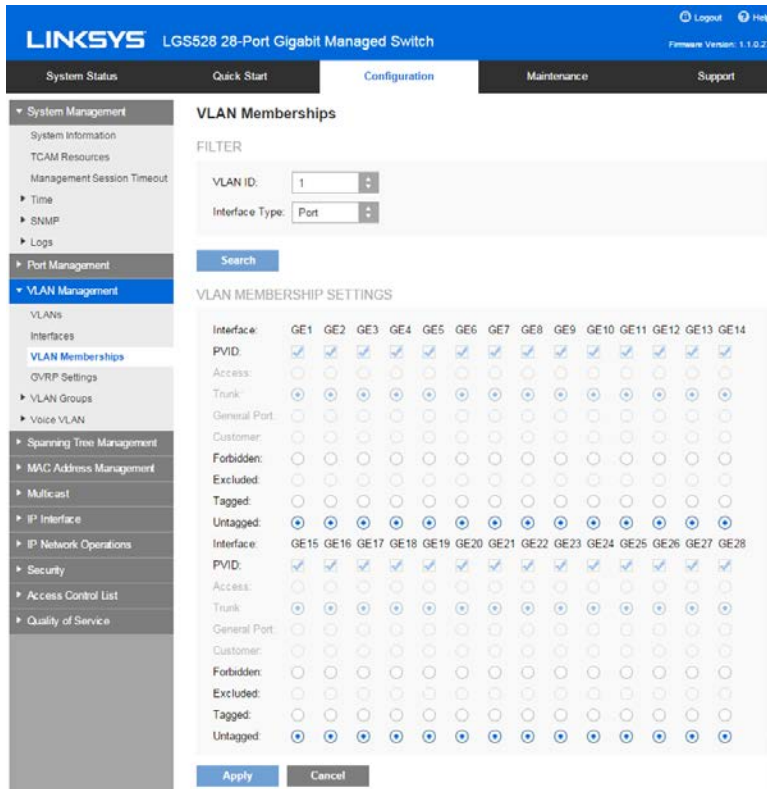
1. Click *Configuration > VLAN Management > Interfaces Settings*.
2. Select an interface type (Port or LAG), and click **Search**.
3. To add a Port or LAG to a VLAN, select it and click **Join VLAN**.

The screenshot shows the 'Join VLAN' configuration window. At the top, it says 'Join VLAN'. Below that, there's a section 'Select Your Interface' with two radio buttons: 'Port' (selected) and 'LAG'. Next to 'Port' is a dropdown menu showing 'GE13', and next to 'LAG' is a dropdown menu showing '1'. Below this is a section 'VLAN Memberships & Settings'. Under 'Mode', it says 'Trunk'. There are two lists for 'VLAN Memberships': one labeled 'VLAN' which is empty, and another labeled 'VLAN Memberships:' which contains '1UP'. Below the lists are 'Tagging' options: 'Forbidden', 'Excluded', 'Tagged' (selected), and 'Untagged'. There is also a 'PVID' checkbox. At the bottom, there's a legend: 'T - Tagged member', 'U - Untagged member', 'I - Internal used', 'P - PVID', 'G - Guest VLAN', 'F - Forbidden member'. At the very bottom are 'Apply' and 'Close' buttons.

4. Enter the following fields:
 - VLAN Mode—Select the interface mode for the VLAN.
The options are:

- 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.
 - General Port—The interface can support all functions as defined in the IEEE 802.1q specification. The interface can be a tagged or untagged member of one or more VLANs.
- Tagging
 - 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.
 - PVID—Port PVID is set to this VLAN. If the interface is in access mode or trunk mode, the device automatically makes the interface an untagged member of the VLAN. If the interface is in general mode, you must manually configure VLAN membership.
5. Click **Apply**. The port is added to the VLAN and the settings 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 packets properly, intermediate VLAN-aware devices that carry VLAN traffic along the path between end nodes must be manually configured.

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 assign a port to one or more VLANs:

1. Click *Configuration > VLAN Management > VLAN Memberships*.
2. Select VLAN ID and interface type (Port or LAG), and click **Search**. Configure the following fields for the interfaces of the selected type:
 - Interface—Port/LAG ID.
 - PVID—Port PVID is set to this VLAN. If the interface is in access mode or trunk mode, the device automatically makes the interface an untagged member of the VLAN. If the interface is in general mode, you must manually configure VLAN membership.
 - Mode—Interface VLAN mode that was selected in the Interface Settings page.
 - Access—Select to make the interface an access interface on this VLAN.
 - Trunk—Select to make the interface a trunk interface on this VLAN.
 - General Port—The interface can support all functions as defined in the IEEE 802.1q specification. The interface can be a tagged or untagged member of one or more VLANs.
 - Customer—Ensure that the customer port and Q in Q are for managed switch only.
 - 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 when the VLAN is newly created. The port can join the VLAN through GVRP registration.
 - Tagged-The interface is a tagged member of the VLAN. This is not relevant for Access ports.
 - Untagged-The interface is an untagged member of the VLAN. Frames of the VLAN are sent untagged to the interface VLAN. This is not relevant for Access ports.
3. Click **Apply**. The settings are modified and written to the Running Configuration file.

GVRP

The screenshot shows the configuration page for a LINKSYS LGS528 28-Port Gigabit Managed Switch. The GVRP Settings are displayed, with the GVRP Status set to 'Enable'. Below this is a table titled 'GVRP Table' with columns for Interface, GVRP State, GVRP VLAN Registration, and Dynamic VLAN Creation. The table lists 28 interfaces (GE1 through GE28). GE13 is selected, and its GVRP State is 'Disabled', GVRP VLAN Registration is 'Enabled', and Dynamic VLAN Creation is 'Enabled'. There is an 'Edit' button at the bottom of the table.

| Interface | GVRP State | GVRP VLAN Registration | Dynamic VLAN Creation |
|-----------|------------|------------------------|-----------------------|
| GE1 | Disabled | Enabled | Enabled |
| GE2 | Disabled | Enabled | Enabled |
| GE3 | Disabled | Enabled | Enabled |
| GE4 | Disabled | Enabled | Enabled |
| GE5 | Disabled | Enabled | Enabled |
| GE6 | Disabled | Enabled | Enabled |
| GE7 | Disabled | Enabled | Enabled |
| GE8 | Disabled | Enabled | Enabled |
| GE9 | Disabled | Enabled | Enabled |
| GE10 | Disabled | Enabled | Enabled |
| GE11 | Disabled | Enabled | Enabled |
| GE12 | Disabled | Enabled | Enabled |
| GE13 | Disabled | Enabled | Enabled |
| GE14 | Disabled | Enabled | Enabled |
| GE15 | Disabled | Enabled | Enabled |
| GE16 | Disabled | Enabled | Enabled |
| GE17 | Disabled | Enabled | Enabled |
| GE18 | Disabled | Enabled | Enabled |
| GE19 | Disabled | Enabled | Enabled |
| GE20 | Disabled | Enabled | Enabled |
| GE21 | Disabled | Enabled | Enabled |
| GE22 | Disabled | Enabled | Enabled |
| GE23 | Disabled | Enabled | Enabled |
| GE24 | Disabled | Enabled | Enabled |
| GE25 | Disabled | Enabled | Enabled |
| GE26 | Disabled | Enabled | Enabled |
| GE27 | Disabled | Enabled | Enabled |
| GE28 | Disabled | Enabled | Enabled |

Adjacent VLAN-aware devices can exchange VLAN information with each other by using Generic VLAN Registration Protocol (GVRP). GVRP is based on the Generic Attribute Registration Protocol (GARP) and propagates VLAN information throughout a bridged network.

Since GVRP requires support for tagging, the port must be configured in Trunk or General mode.

When a port joins a VLAN by using GVRP, it is added to the VLAN as a dynamic member, unless this was expressly forbidden in the Port VLAN Membership page. If the VLAN does not exist, it is dynamically created when Dynamic VLAN creation is enabled for this port (in the GVRP Settings page).

GVRP must be activated globally as well as on each port. When it is activated, it transmits and receives GARP Packet Data Units (GPDUs). VLANs that are defined but not active are not propagated. To propagate the VLAN, it must be up on at least one port.

By default, GVRP is disabled globally and on ports.

To define GVRP settings for an interface:

1. Click *Configuration > VLAN Management > GVRP Settings*.
2. Select GVRP Global Status to enable GVRP globally.
3. Click **Apply** to set the global GVRP status.
4. Select an interface type (Port or LAG), and click **Search** to display all interfaces of that type.
5. To define GVRP settings for a port, select it, and click **Edit**.
6. Enter the values for the following fields:
 - Interface—Select the interface (Port or LAG) to be edited.
 - GVRP State—Select to enable GVRP on this interface.
 - GVRP VLAN Registration—Select to enable VLAN Registration using GVRP on this interface.
 - Dynamic VLAN Creation—Select to enable Dynamic VLAN Creation on this interface.
7. Click **Close**. GVRP settings are modified, and written to the Running Configuration file.

VLAN Groups

This section describes how to configure VLAN groups. It describes the following processes:

- MAC-Based Groups
- Protocol-based VLAN

VLAN groups classify packets into VLANs based on their MAC addresses or protocol ID.

VLAN groups can be used to separate traffic into different VLANs for security and/or load balancing.

If several classifications schemes are defined, packets are assigned to a VLAN in the following order:

- TAG—If the packet is tagged, the VLAN is taken from the tag.
- MAC-Based VLAN—If a MAC-based VLAN has been defined, the VLAN is taken from the source MAC-to-VLAN mapping of the ingress interface.
- Protocol-Based VLAN—If a protocol-based VLAN has been defined, the VLAN is taken from the (Ethernet type) protocol-to-VLAN mapping of the ingress interface.
- PVID—VLAN is taken from the port default VLAN ID.

MAC-based Groups

MAC-based VLAN classification enable packets to be classified according to their source MAC address. You can then define MAC-to-VLAN mapping per interface.

You can define several MAC-based VLAN groups, which each group containing different MAC addresses.

These MAC-based groups can be assigned to specific ports/LAGs. MAC-based VLAN groups cannot contain overlapping ranges of MAC addresses on the same port.

The following table describes the availability of MAC-based VLAN groups in various SKUs:

Table 1 MAC-Based VLAN Group Availability

| SKU | System Mode | MAC-based VLAN Groups Supported |
|---------|-------------|---------------------------------|
| Smart | Layer 2 | Yes |
| | Layer 3 | No |
| Managed | Layer 2 | Yes |
| | Layer 3 | No |

To define a MAC-based VLAN group:

Assign a MAC address to a VLAN group ID (using the MAC-Based Groups page).

For each required interface:

- Assign the VLAN group to a VLAN (using the Mac-Based VLAN page). The interfaces must be in General mode.
- If the interface does not belong to the VLAN, manually assign it to the VLAN using the VLAN Membership page.

To assign a MAC address to a VLAN Group:

1. Click *Configuration > VLAN Management > VLAN Groups > MAC-Based Group*.
2. Click **Add**.

Add MAC-Based Group

New MAC-Based Group

Group ID: (1 - 2147483647)

Group Settings

MAC Address:

Prefix Mask: Host (48) Length (9 - 48)

3. Enter the values for the following fields:

- Group ID—Enter a user-created VLAN group ID number.
- MAC Address—Enter a MAC address to be assigned to a VLAN group.

Note—*This MAC address cannot be assigned to any other VLAN group.*

- Prefix Mask—Enter one of the following:
- Host—Source host of the MAC address
- Length—Prefix of the MAC address

4. Click **Apply**. The MAC address is assigned to a VLAN group.

To assign a MAC-based VLAN group to a VLAN on an interface:

1. Click *Configuration > VLAN Management > VLAN Groups > MAC-Based VLAN*.
2. Click **Add**.

Add MAC-Based VLAN

Select Your Interface

Interface: Port LAG

MAC-Based VLAN Settings

Group ID: VLAN ID: (1-4094)

3. Enter the values for the following fields:
 - Interface—Enter a general interface (port/LAG) through which traffic is received.
 - Group ID—Select a VLAN group, defined in the MAC-Based Groups page.
 - VLAN ID—Select the VLAN to which traffic from the VLAN group is forwarded.
4. Click **Apply** to set the mapping of the VLAN group to the VLAN. This mapping does not bind the interface dynamically to the VLAN. The interface must be manually added to the VLAN.

Protocol-based VLAN

Groups of protocols can be defined and then bound to a port. After the protocol group is bound to a port, every packet originating from a protocol in the group is assigned the VLAN that is configured in the Protocol-Based Groups page.

To define a protocol-based VLAN group:

- Define a protocol group (using the Protocol-Based Groups page).
- For each required interface, assign the protocol group to a VLAN (using Protocol-based VLAN page). The interfaces must be in General mode and cannot have a Dynamic VLAN (DVA) assigned to it.

To define a set of protocols:

1. Click *Configuration > VLAN Management > Protocol-Based Group*.

The Protocol-Based Groups Page contains the following fields:

- Group ID—Displays the protocol group ID to which the interface is added.
- Encapsulation—Displays the protocol on which the VLAN group is based.
- Protocol Value / DSAP-SSAP—Displays the protocol value in hex.

2. Click the **Add**.

Add Protocol-Based Group

New Protocol-Based Group

Group ID: (1-2147483647)

Group Settings

Encapsulation: Ethernet V2 LLC-SNAP LLC

Ethernet Type: ▾

Protocol Value: (0x0600-0xFFFF)

Apply
Close

3. Enter the following fields:

- Group ID—Enter a protocol group ID.
- Encapsulation—Protocol Packet type. The following options are available:
- Ethernet V2—If this is selected, select the Ethernet Type.
- LLC-SNAP (rfc1042)—If this is selected, enter the Protocol Value.
- LLC—If this is selected, select the DSAP-SSAP Values.
- Ethernet Type—Select the Ethernet type for Ethernet V2 encapsulation. This is the two-octet field in the Ethernet frame used to indicate which protocol is encapsulated in the payload of the Ethernet packet) for the VLAN group
- Protocol Value—Enter the protocol for LLC-SNAP (rfc 1042) encapsulation.
- DSAP-SSAP—Enter these values for LLC encapsulation.

4. Click **Apply**. The Protocol Group is added, and written to the Running Configuration file.

To map a protocol group to a port, the port must be in General mode and not have DVA configured on it (see Interface Settings).

Several groups can be bound to a single port, with each port being associated to its own VLAN.

It is possible to map several groups to a single VLAN as well.

To map the protocol port to a VLAN:

1. Click *Configuration > VLAN Management > Protocol-Based VLAN*.
2. To associate an interface with a protocol-based group and VLAN, click **Add**.

The screenshot shows a configuration window titled "Add Protocol-Based VLAN". It has two main sections: "Select Your Interface" and "Protocol-Based VLAN Settings".

Select Your Interface: This section contains two radio buttons. The "Port" radio button is selected. Next to it is a dropdown menu showing "GE1". The "LAG" radio button is unselected. Next to it is a dropdown menu showing "1".

Protocol-Based VLAN Settings: This section contains two input fields. The "Group ID" field is empty. The "VLAN ID" field contains the number "1". To the right of the "VLAN ID" field is the text "(1-4094)".

At the bottom of the dialog, there are two buttons: "Apply" (in blue) and "Close" (in grey).

3. Enter the following fields.

- Interface—Port or LAG number assigned to VLAN according to protocol- based group.
- Group ID—Protocol group ID.
- VLAN ID—Attaches the interface to a user-defined VLAN ID.

4. Click **Apply**. The protocol ports are mapped to VLANs, and written to the Running Configuration file.

Voice VLAN

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

Overview

Auto Voice VLAN

The device supports the Telephony OUI (Organization Unique Identifier) voice VLAN mode. The two modes affect how voice VLAN and/or voice VLAN port memberships are configured.

In Telephony OUI mode, the voice VLAN must be a manually-configured VLAN, and cannot be the default VLAN.

When the device is in Telephony OUI mode and 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 one of the configured telephony OUIs. An OUI is the first three bytes of an Ethernet MAC address. For more information about Telephony OUI, see Telephony OUI.

Voice End-Points

To have a voice VLAN work properly, the 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:

- 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.
- A phone/endpoint may obtain the voice VLAN information from CDP and LLDP-MED advertisements it receives from their neighbor voice systems and switches.

You can create a network policy manually or enable the device to automatically generate a network policy, based on a voice VLAN configuration.

The device expects the attaching voice devices to send voice VLAN, tagged packets. On ports where the voice VLAN is the native VLAN or that is configured with auto voice VLAN by Telephony OUI, voice VLAN untagged packets are possible.

Voice VLAN QoS

The device can advertise the CoS/802.1p and DSCP settings of the voice VLAN by using LLDP-MED Network policies. You can create your network policy manually or enable the device to automatically generate the network policy based on your voice VLAN configuration. MED-supported devices must send their voice traffic with the same CoS/802.1p and DSCP values, as received with the LLDP-MED response.

You can disable the automatic update between Voice VLAN and LLDP-MED and use his own network policies.

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.

In Auto Voice VLAN, you can override the value of the voice streams using advanced QoS. 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.

In addition the following constraints are applicable for Telephony OUI:

- The Voice VLAN cannot be VLAN1 (the default 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 General or 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).

Feature Configuration

The screenshot shows the configuration page for the Linksys LGS528 28-Port Gigabit Managed Switch. The page is titled "Feature Configuration" and is part of the "Voice VLAN" section under "VLAN Management". The configuration options are as follows:

- VOICE VLAN:**
 - Voice VLAN ID: 1 (range 1-4094)
 - CoS/802.1p: 6
- TELEPHONY OUI:**
 - Telephone OUI Voice VLAN: Enabled
 - Remark CoS/802.1p: Enabled
 - Aging Time: 1 days, 0 hours, 0 min (1min - 30days)

Buttons for "Apply" and "Cancel" are present. Below the configuration is a "Telephony OUI Table" with the following entries:

| Telephone OUI | Description |
|-----------------------------------|-------------------------------|
| <input type="checkbox"/> 00-01-e3 | Siemens_AG_phone_____ |
| <input type="checkbox"/> 00-03-6b | Cisco_phone_____ |
| <input type="checkbox"/> 00-09-6e | Avaya_____ |
| <input type="checkbox"/> 00-0f-e2 | H3C_Aolynk_____ |
| <input type="checkbox"/> 00-60-b9 | Philips_and_NEC_AG_phone_____ |
| <input type="checkbox"/> 00-d0-1e | Pingtel_phone_____ |
| <input type="checkbox"/> 00-e0-75 | Polycom/Ventel_phone_____ |
| <input type="checkbox"/> 00-e0-bb | 3Com_phone_____ |

Buttons for "Add", "Delete", and "Restore" are located at the bottom of the table.

To configure Auto Voice VLAN:

1. Click *Configuration* > *VLAN Management* > *Voice VLAN* > *Feature Configuration*.
2. Enter the following to configure Voice VLAN:
 - Voice VLAN ID—Enter the identifier of the current voice VLAN
 - CoS/802.1p—Select the CoS/802.1p value to be used by the LLDP-MED as a voice network policy.
3. Enter the following to configure telephone OUI voice VLAN:
 - Telephone OUI Voice VLAN—Check to enable automatically adding ports to voice VLAN when OUI packets are received.
 - Remark CoS/802.1p—Select the enable remarking packets with the CoS/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.
4. Click **Apply** to save the settings to the Running Configuration file.

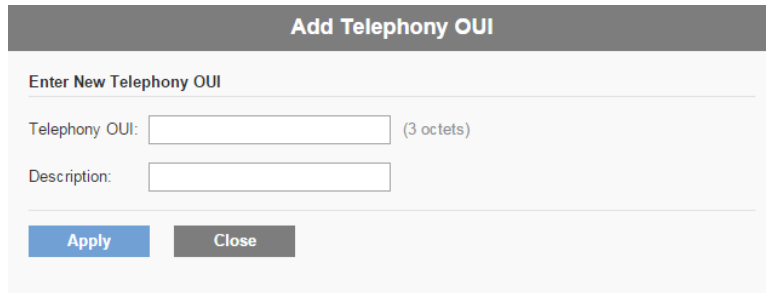
Note—Refer to [Configuration > Port Management > Discovery - LLDP > LLDP MED Network Policy](#) to enable automatic generation of network policy for voice.

To view or add a new OUI:

1. Click [Configuration > VLAN Management > Voice VLAN > Feature Configuration](#).
The Telephony OUI page displays the configured OUIs.

Note—Click [Restore](#) to delete all of the user-created OUIs, and leave only the default OUIs in the table. The OUI information may not be accurate until the restoration is completed. This may take several seconds. After several seconds have passed, refresh the page by exiting it and re-entering it.

2. Click **Add** to add a new OUI.



The screenshot shows a modal dialog box titled "Add Telephony OUI". Below the title bar, it says "Enter New Telephony OUI". There are two input fields: "Telephony OUI:" followed by a text box and "(3 octets)" to its right, and "Description:" followed by another text box. At the bottom of the dialog, there are two buttons: "Apply" (highlighted in blue) and "Close".

3. Enter the fields:
 - a. Telephony OUI—First six digits of the MAC address that are reserved for OUIs.
 - b. Description—User-assigned OUI description.
4. Click **Apply**. The OUI is added to the Telephony OUI Table.

Telephone OUI Interfaces

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 Interface 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:

1. Click *Configuration > VLAN Management > Voice VLAN > Telephony OUI Interfaces*.
Voice VLAN OUI parameters for all interfaces are displayed.
2. To configure an interface to be a candidate port of the telephony OUI-based voice VLAN, click **Edit**.
3. Enter the values for the following fields:
 - **Interface**—Select an interface.
 - **Telephony OUI VLAN**—If enabled, 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.
4. Click **Apply**. The OUI Interface is added.

Chapter 7 - Spanning Tree

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](#)
- [MSTP Properties](#)
- [VLAN to MSTP](#)
- [MSTP Instance Status](#)
- [MSTP Instance Interface](#)

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.

- Multiple STP (MSTP) – Classic STP and Rapid STP detect Layer 2 loops, and attempt to mitigate them by preventing the involved port from transmitting traffic. Since loops exist on a per-Layer 2-domain basis, there can be a loop in VLAN A and no loop in VLAN B. If both VLANs are on Port X, Classic STP and Rapid STP will mitigate the loop by stopping traffic on the entire port, including VLAN B traffic.

MSTP solves this problem by enabling several STP instances, so that it is possible to detect and mitigate loops separately in each instance. By associating instances to VLANs, each instance is associated with the Layer 2 domain on which it performs loop detection and mitigation. This enables a port to be stopped in one instance, such as traffic from VLAN A that is causing a loop, while traffic can remain active in another domain where no loop was seen, such as on VLAN B.

Spanning Tree

The Spanning Tree page contains parameters for enabling STP, RSTP, or MSTP.

To set the STP status and global settings:

- Click *Configuration > Spanning Tree Management > Spanning Tree*.

The screenshot shows the configuration page for Spanning Tree on a Linksys switch. The page is titled "Spanning Tree" and includes a navigation menu on the left. The main content area is divided into "SETTINGS" and "STATUS" sections.

SETTINGS

Global

- Spanning Tree: Enable
- Spanning Tree Mode: Classic STP Rapid STP Multiple STP
- Path Cost Default Values: Short Long
- BPDU Handling: Filtering Flooding

Bridge Configurations

- Priority: (0-61440) Hello Time: sec (1-10)
- Maximum Age: sec (6-40) Forward Delay: sec (4-30)

STATUS

Designated Root

| | | | |
|-------------------------|-------------------------|-----------------------|-------------------------|
| Bridge ID: | 32768-b4-75-0e-7c-5b-7a | Root Bridge ID: | 32768-b4-75-0e-2d-98-9e |
| Root Port: | GE13 | Root Path Cost: | 20000 |
| Topology Changes Count: | 1 | Last Topology Change: | 0D/3H/30M/55S |

Buttons:

- Enter the parameters.

- Global Settings:
 - Spanning Tree—Select to enable on the device.
 - Spanning Tree Mode—Select an STP mode - Classic STP, Rapid STP or Multiple 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.
 - o Short—Specifies the range 1 through 65,535 for port path costs.
 - o 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.
 - o Filtering—Filters BPDU packets when Spanning Tree is disabled on an interface.
 - o Flooding—Floods BPDU packets when Spanning Tree is disabled on an interface.
- Bridge Settings:
 - Priority—Sets the bridge priority value. After exchanging BPDUs, the device with the lowest priority becomes the Root Bridge. In the case that all bridges use the same priority, then their MAC addresses are used to determine the Root Bridge. The bridge priority value is provided in increments of 4096. For example, 4096, 8192, 12288, and so on.
 - Hello Time—Set the interval (in seconds) that a Root Bridge waits between configuration messages.
 - Max 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 packet.
- Designated Root:
 - Bridge ID—The bridge priority concatenated with the MAC address of the device.
 - Root Bridge ID—The Root Bridge priority concatenated with 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 Counts—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.

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 versions of the STP protocol.

To configure STP on an interface:

1. Click *Configuration > Spanning Tree Management > STP Interfaces*.
2. Select an interface and click **Edit**.

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.
 - Root Guard—Click to enable.
 - BPDU Guard—Click to enable.
 - 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.
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 Interface Settings 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 or MSTP.

To configure RSTPs:

1. Click *Configuration > Spanning Tree Management > Spanning Tree*.
2. Enable Rapid STP.
3. Click *Configuration > Spanning Tree Management > Spanning Tree > RSTP Interfaces*.
4. Select an interface type, choose a port and click **Edit**.

Edit RSTP Interface

Select Your Interface

Interface: Port GE13 LAG 1

Interface Settings

Point to Point Mode: Enable Disable Auto

Edge Port Mode: Enable Disable Auto

STP Mode: RSTP Point to Point Status: Enabled

Port Role: Root Port Status: Forwarding

Apply Close

5. 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.
6. Click **Apply**. The Running Configuration file is updated.

MSTP Properties

Multiple Spanning Tree Protocol (MSTP) is used to separate the STP port state between various domains (on different VLANs). For example, while port A is blocked in one STP instance due to a loop on VLAN A, the same port can be placed in the Forwarding State in another STP instance. The MSTP Properties page enables you to define the global MSTP settings.

To configure MSTP:

1. Set the STP Operation Mode to Multiple STP as described in the Spanning Tree page.
2. Define MSTP instances. Each MSTP instance calculates and builds a loop free topology to bridge packets from the VLANs that map to the instance. Refer to the MSTP Properties page.
3. Decide which MSTP instance will be active in what VLAN, and associate these MSTP instances to VLAN(s) accordingly.
4. Configure the MSTP attributes by the following pages:
 - MSTP Properties
 - MSTP Instance Status
 - MSTP Instance Interface

MSTP Interfaces

The global MSTP configures a separate Spanning Tree for each VLAN group and blocks all but one of the possible alternate paths within each spanning tree instance. MSTP enables formation of MST regions that can run multiple MST instances (MSTI). Multiple regions and other STP bridges are interconnected using one single common spanning tree (CST).

MSTP is fully compatible with RSTP bridges, in that an MSTP BPDU can be interpreted by an RSTP bridge as an RSTP BPDU. This not only enables compatibility with RSTP bridges without configuration changes, but also causes any RSTP bridges outside of an MSTP region to see the region as a single RSTP bridge, regardless of the number of MSTP bridges inside the region itself. Up to seven MST instances can be defined on the managed switches in addition to instance zero.

VLAN to MSTP

For two or more switches to be in the same MST region, they must have the same VLANs to MST instance mapping, the same configuration revision number, and the same region name.

Switches intended to be in the same MST region are never separated by switches from another MST region. If they are separated, the region becomes two separate regions.

The VLAN to MSTP instance mapping is done in the MSTP Properties page. Each VLAN can be mapped to a MSTP instance. For devices to be in the same region, they must have the same mapping of VLANs to MSTP instances.

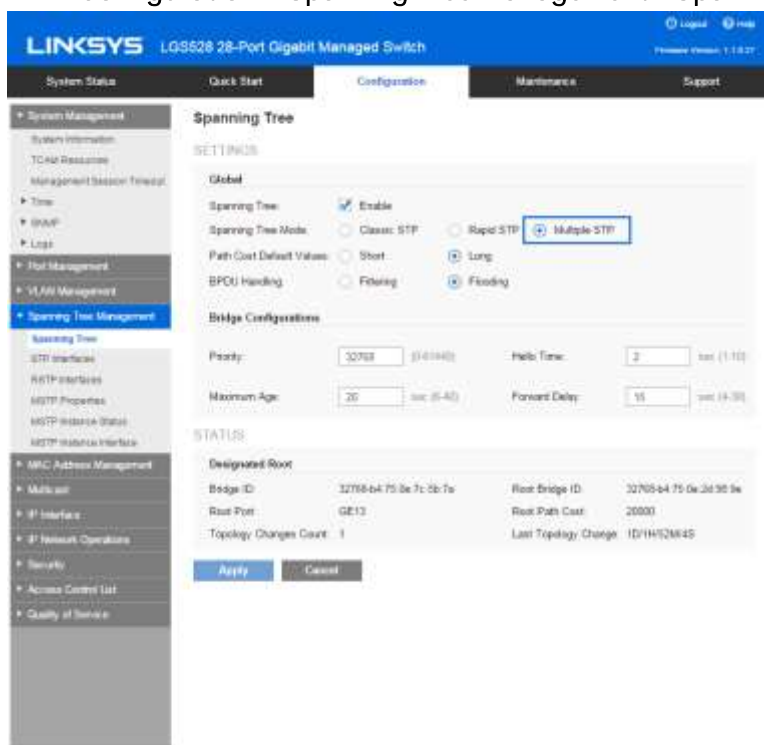
Configuration on this page (and all of the MSTP pages) applies if the system STP mode is MSTP.

Note—The same MSTI can be mapped to more than one VLAN, but each VLAN can only have one MST Instance attached to it.

For those VLANs that are not explicitly mapped to one of the MST instances, the device automatically maps them to the CIST (Core and Internal Spanning Tree) instance. The CIST instance is MST instance 0.

To configure MSTP:

1. Click **Configuration > Spanning Tree Management > Spanning Tree > Enable MSTP.**



2. Click **Apply.**

3. Click Configuration > Spanning Tree Management> MSTP Properties.

Note—IST Master, which displays the regions master, is a display only field.

The screenshot shows the Linksys web interface for a LG5528 28-Port Gigabit Managed Switch. The main content area is titled 'MSTP Properties' and is divided into 'SETTINGS' and 'STATUS' sections. The 'SETTINGS' section includes fields for 'Region Name' (set to '34.75.26.7c.5b.7a'), 'Revision' (set to '8'), and 'Maximum Hops' (set to '20'). The 'STATUS' section shows the 'IST Master' as '32768.64.75.26.7c.5b.7a'. Below the settings are 'Apply' and 'Cancel' buttons. At the bottom, there is an 'MSTP Instance Table' with columns for 'MSTP Instance ID', 'Bridge Priority', and 'VLANs'. The table contains seven rows, each with a radio button, an instance ID (1-7), a bridge priority of 32768, and a VLANs field.

| MSTP Instance ID | Bridge Priority | VLANs |
|-------------------------|-----------------|-------|
| <input type="radio"/> 1 | 32768 | |
| <input type="radio"/> 2 | 32768 | |
| <input type="radio"/> 3 | 32768 | |
| <input type="radio"/> 4 | 32768 | |
| <input type="radio"/> 5 | 32768 | |
| <input type="radio"/> 6 | 32768 | |
| <input type="radio"/> 7 | 32768 | |

4. Enter the parameters.
 - Region Name—Define an MSTP region name.
 - Revision—Define an unsigned 16-bit number that identifies the revision of the current MST configuration. The field range is from 0 to 65535.
 - Maximum Hops—Set the total number of hops that occur in a specific region before the BPDU is discarded. Once the BPDU is discarded, the port information is aged out. The field range is from 1 to 40.
5. Click Apply. The MSTP properties are defined, and the Running Configuration file is updated.

To configure an MSTP instance:

1. Click *Configuration > Spanning Tree Management > MSTP Properties*.
2. Select the MST instance and click **Edit**.

Edit STP Interface

Select Your Interface

Interface: Port LAG

Interface Settings

STP: Enable

BPDU Handling: Use Global Settings Filtering Flooding

Path Cost: Use Default User Defined (1-200000000)

Priority:

Root Guard: Enable BPDU Guard: Enable

Port State: Forwarding Designated Bridge ID: 32768-b4-75:0e:2d:98:9e

Designated Port ID: 128-53 Designated Cost: 0

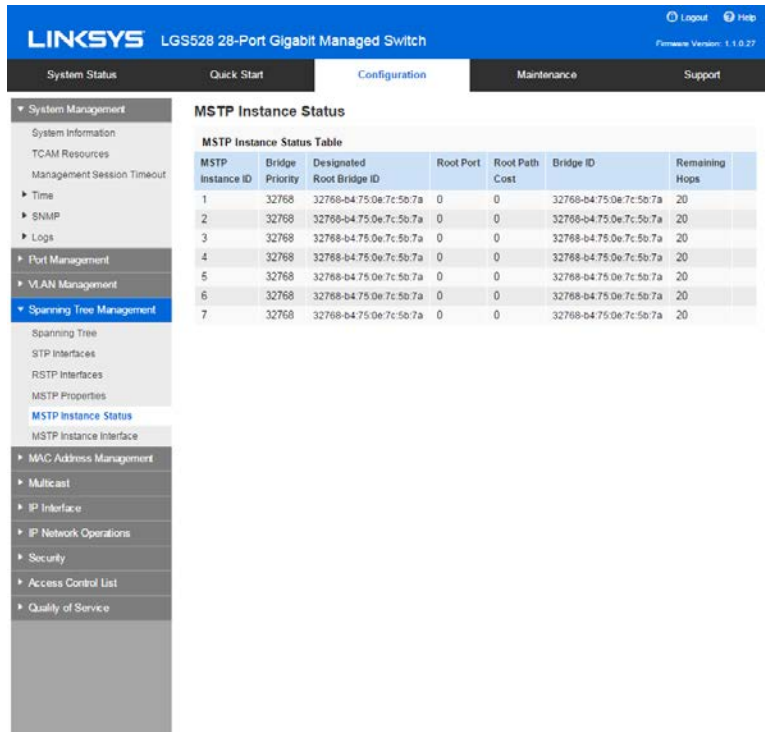
3. Enter the parameters.
 - MST Instance ID—Select an MST instance to be displayed and defined.
 - Bridge Priority—Set the priority of this bridge for the selected MST instance.
 - Action—Select either *Add VLAN* or *Remove VLAN*.
 - VLANs—Displays the VLANs mapped to the selected instance. The default mapping is that all VLANs are mapped to the common and internal spanning tree (CIST) instance 0.
4. Click **Apply**. The MSTP properties are defined, and the Running Configuration file is updated.

MSTP Instance Status

The MSTP Instance Status page enables you to configure and view parameters per MST instance. This is the per-instance equivalent to the Spanning Tree page.

To enter MSTP instance settings:

Click *Configuration > Spanning Tree Management > MSTP Instance Status*.



The screenshot shows the Linksys LGS528 28-Port Gigabit Managed Switch web interface. The top navigation bar includes 'System Status', 'Quick Start', 'Configuration', 'Maintenance', and 'Support'. The left sidebar lists various system management options, with 'Spanning Tree Management' selected. The main content area displays the 'MSTP Instance Status' page, which includes a table titled 'MSTP Instance Status Table'.

| MSTP Instance ID | Bridge Priority | Designated Root Bridge ID | Root Port | Root Path Cost | Bridge ID | Remaining Hops |
|------------------|-----------------|---------------------------|-----------|----------------|-------------------------|----------------|
| 1 | 32768 | 32768-b4.75.0e.7c.5b.7a | 0 | 0 | 32768-b4.75.0e.7c.5b.7a | 20 |
| 2 | 32768 | 32768-b4.75.0e.7c.5b.7a | 0 | 0 | 32768-b4.75.0e.7c.5b.7a | 20 |
| 3 | 32768 | 32768-b4.75.0e.7c.5b.7a | 0 | 0 | 32768-b4.75.0e.7c.5b.7a | 20 |
| 4 | 32768 | 32768-b4.75.0e.7c.5b.7a | 0 | 0 | 32768-b4.75.0e.7c.5b.7a | 20 |
| 5 | 32768 | 32768-b4.75.0e.7c.5b.7a | 0 | 0 | 32768-b4.75.0e.7c.5b.7a | 20 |
| 6 | 32768 | 32768-b4.75.0e.7c.5b.7a | 0 | 0 | 32768-b4.75.0e.7c.5b.7a | 20 |
| 7 | 32768 | 32768-b4.75.0e.7c.5b.7a | 0 | 0 | 32768-b4.75.0e.7c.5b.7a | 20 |

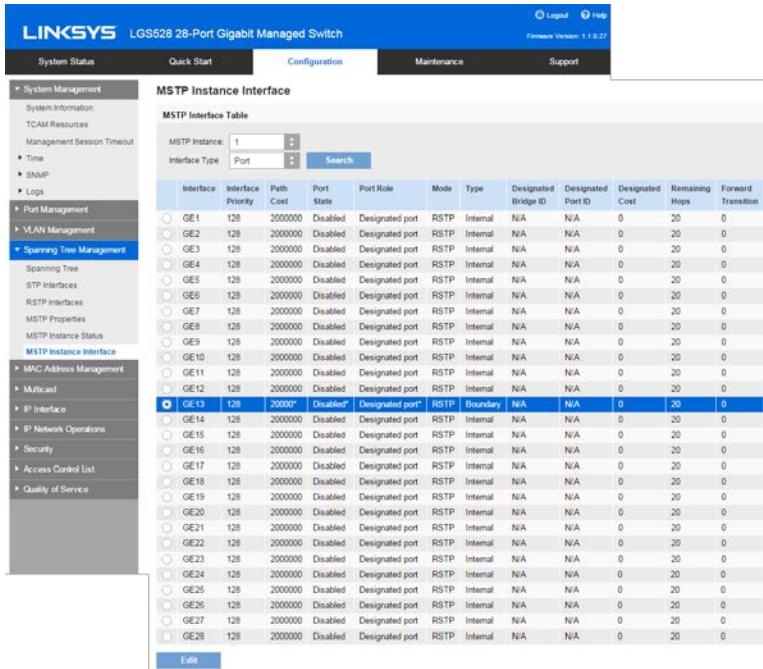
- Instance ID—Select an MST instance to be displayed and defined.
- Bridge Priority—Set the priority of this bridge for the selected MST instance.
- Designated Root Bridge ID—Displays the priority and MAC address of the Root Bridge for the MST instance.
- Root Port—Displays the root port of the selected instance.
- Root Path Cost—Displays the root path cost of the selected instance.
- Bridge ID—Displays the bridge priority and the MAC address of this device for the selected instance.
- Remaining Hops—Displays the number of hops remaining to the next destination.

MSTP Instance Interface

The MSTP Instance Interface page enables you to configure the port MSTP settings for every MST instance, and to view information that has currently been learned by the protocol, such as the designated bridge per MST instance.

To configure the ports in an MST instance:

1. Click *Configuration > Spanning Tree Management > MSTP Instance Interface*.



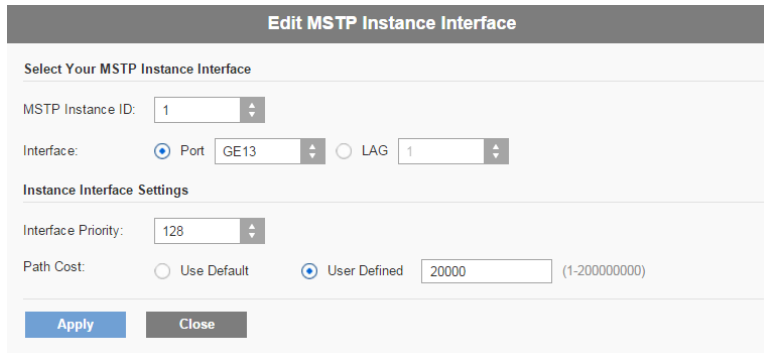
2. Enter the parameters.
 - MSTP Instance—Select the MSTP instance to be configured.
 - Interface Type—Select whether to display the list of ports or LAGs.
3. Click **Search**. The following MSTP parameters for the interfaces on the instance are displayed:
 - Interface—Select the interface for which the MSTI settings are to be defined.
 - Interface Priority—Set the port priority for the specified interface and MST instance.
 - Path Cost—Enter the port contribution to the root path cost in the User Defined textbox or select Use Default to use the default value.
 - Port State—Displays the MSTP status of the specific port on a specific MST instance. The parameters are defined as:
 - Disabled—STP is currently disabled.
 - Blocking—The port on this instance is currently blocked, and cannot forward traffic (with the exception of BPDU data) or learn MAC addresses.
 - Listening—The port on this instance is in Listening mode. The port cannot forward traffic, and cannot learn MAC addresses.

- Learning—The port on this instance is in Learning mode. The port cannot forward traffic, but it can learn new MAC addresses.
- Forwarding—The port on this instance is in Forwarding mode. The port can forward traffic and learn new MAC addresses.
- Boundary—The port on this instance is a boundary port. It inherits its state from instance 0 and can be viewed on the STP Interface Settings page.
- Port Role—Displays the port or LAG role, per port or LAG per instance, assigned by the MSTP algorithm to provide STP paths:
 - Root—Forwarding packets through this interface provides the lowest cost path for forwarding packets to the root device.
 - Designated—The interface through which the bridge is connected to the LAN, which provides the lowest root path cost from the LAN to the Root Bridge for the MST instance.
 - Alternate—The interface provides an alternate path to the root device from the root interface.
 - Backup—The interface provides a backup path to the designated port path toward the Spanning Tree leaves. Backup ports occur when two ports are connected in a loop by a point-to-point link. Backup ports also occur when a LAN has two or more established connections to a shared segment.
 - Disabled—The interface does not participate in the Spanning Tree.
 - Boundary—The port on this instance is a boundary port. It inherits its state from instance 0 and can be viewed on the STP Interface Settings page.
- Mode—Displays the current interface Spanning Tree mode.
 - If the link partner is using MSTP or RSTP, the displayed port mode is RSTP.
 - If the link partner is using STP, the displayed port mode is STP.
- Type—Displays the MST type of the port.
 - Boundary—A Boundary port attaches MST bridges to a LAN in a remote region. If the port is a boundary port, it also indicates whether the device on the other side of the link is working in RSTP or STP mode.
 - Internal—The port is an internal port.
- Designated Bridge ID—Displays the ID number of the bridge that connects the link or shared LAN to the root.
- Designated Port ID—Displays the Port ID number on the designated bridge that connects the link or the shared LAN to the root.

- 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.
- Remain Hops—Displays the hops remaining to the next destination.
- Forward Transition—Displays the number of times the port has changed from the Forwarding state to the Blocking state.

4. Select an interface, and click **Edit**.

5. Enter the parameters.



The screenshot shows a configuration window titled "Edit MSTP Instance Interface". It is divided into two main sections: "Select Your MSTP Instance Interface" and "Instance Interface Settings".

Select Your MSTP Instance Interface:

- MSTP Instance ID: 1
- Interface: Port GE13 LAG 1

Instance Interface Settings:

- Interface Priority: 128
- Path Cost: Use Default User Defined 20000 (1-200000000)

At the bottom, there are two buttons: "Apply" and "Close".

6. Click **Apply**. The Running Configuration file is updated.

Chapter 8 - MAC Address Management

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

- [Dynamic MAC Addresses](#)
- [Static MAC Addresses](#)
- [Reserved MAC Addresses](#)

There are two types of MAC addresses—static and dynamic. Depending on their type, MAC addresses are either stored in the Static Address table or in the Dynamic Address table, along with VLAN and port information.

Static addresses are configured by the user, and therefore, they do not expire.

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 static or 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 16K static and dynamic MAC addresses.

Dynamic MAC Addresses

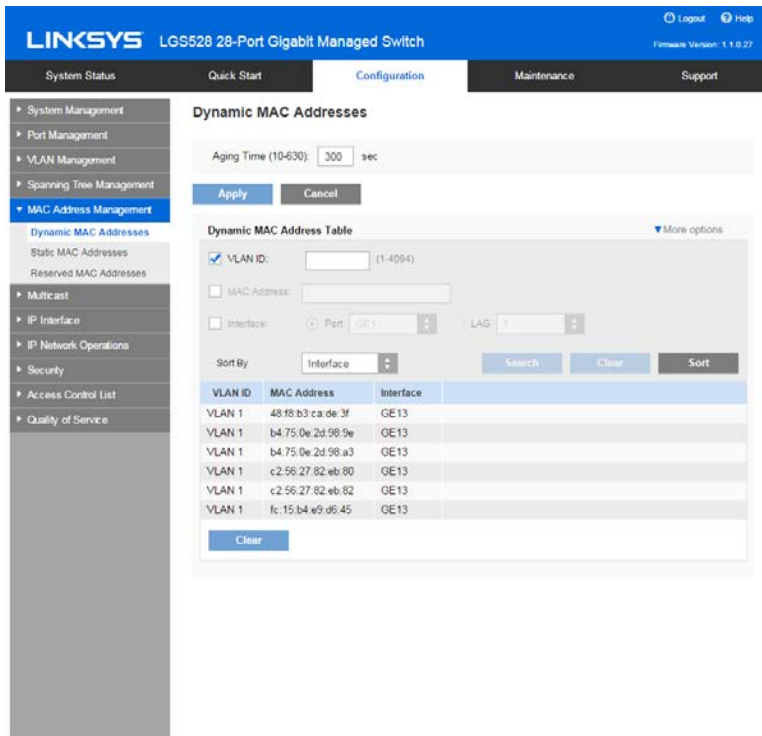
The Dynamic Address Table (bridging table) contains the MAC addresses acquired by monitoring the source addresses of frames entering the device.

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 of time known as the aging time.

Configuring Dynamic MAC Address Aging Time

To configure the aging time for dynamic addresses:

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



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.
3. Click **Apply**. The aging time is updated.
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.
5. Click **Search**. The Dynamic MAC Address Table is queried and the results are displayed.
 6. To delete all of the dynamic MAC addresses, click **Clear**.

Static MAC Addresses

Static MAC addresses are assigned to a specific physical interface and VLAN on the device. If that address is detected on another interface, it is ignored, and is not written to the address table.

To define a static address:

1. Click *Configuration > MAC Address Management > Static MAC Addresses*.
The Static Addresses page contains the currently defined static addresses.
2. Click **Add**.

3. Enter the parameters.
 - VLAN ID—Select the VLAN ID for the port.
 - MAC Address—Enter the interface MAC address.
 - Interface—Select an interface (unit/slot, port, or LAG) for the entry.
 - Status—Select how the entry is treated. The options are:
 - Permanent—The system never removes this MAC address. If the static MAC address is saved in the Startup Configuration, it is retained after rebooting.

- Delete on reset—The static MAC address is deleted when the device is reset.
- Delete on timeout—The MAC address is deleted when aging occurs.
- Secure—The MAC address is secure when the interface is in classic locked mode (see Configuring Port Security).

4. Click **Apply**. A new entry appears in the table.

Reserved MAC Addresses

When the device receives a frame with a destination MAC address that belongs to a reserved range (per the IEEE standard), the frame can be discarded or bridged. The entry in the Reserved MAC Address Table can either specify the reserved MAC address or the reserved MAC address and a frame type:

To add an entry for a reserved MAC address:

1. Click *Configuration > MAC Address Tables > Reserved MAC Addresses*.

The screenshot shows the Linksys web interface for the LGS528 28-Port Gigabit Managed Switch. The page is titled "Reserved MAC Addresses" and is part of the "MAC Address Management" section. The interface includes a navigation menu on the left and a main content area. The main content area displays a table with the following columns: "MAC Address", "Frame Type", "Protocol", and "Action". Below the table, it indicates "0 results found." and provides buttons for "Add", "Edit", and "Delete".

2. Click **Add**.

Add Reserved MAC Address

Enter New MAC Address

MAC Address:

MAC Address Settings

Frame Type: Ethernet II LLC LLC-SNAP All

Action: Bridge Discard

3. Enter the values for the following fields:

- MAC Address—Select the MAC address to be reserved.
- Frame Type—Select a frame type based on the following criteria:
 - Ethernet II—Applies to Ethernet II packets with the specific MAC address and ethertype.
 - LLC—Applies to Logical Link Control (LLC) packets with the specific MAC address and DSAP-SSAP.
 - LLC-SNAP—Applies to Logical Link Control/Sub-Network Access Protocol (LLC-SNAP) packets with the specific MAC address.
 - All—Applies to all packets with the specific MAC address and protocol.
- Action—Select one of the following actions to be taken upon receiving a packet that matches the selected criteria:
 - Bridge—Forward the packet to all VLAN members.
 - Discard—Drop the packet.

4. Click **Apply**. A new MAC address is reserved.

Chapter 9 - Multicast

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

- [Overview](#)
- [Feature Configuration](#)
- [IGMP Snooping](#)
- [MLD Snooping](#)
- [Multicast Router Ports](#)
- [Forward All](#)
- [Unregistered Multicast](#)
- [IGMP/MLD 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 (one of these options can be configured per VLAN):

- 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.

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 or IPv6 group address. If a Multicast client can receive Multicast traffic from any source of a specific Multicast group, this is saved as (*,G).

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).

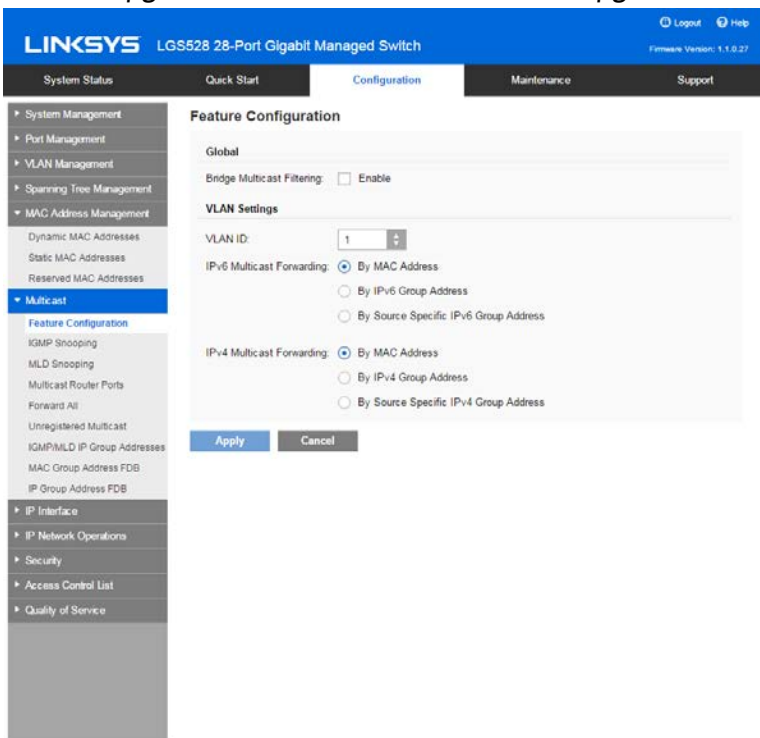
By selecting the forwarding mode, you can define the method used by hardware to identify Multicast flow by one of the following options: MAC Group Address, IP Group Address, or Source Specific IP Group Address.

(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:

1. Click *Configuration > Multicast > Feature Configuration*.



2. Enter the global parameters:
 - Bridge Multicast Filtering—Select to enable filtering of Multicast addresses.
 - VLAN Settings:

- VLAN ID—Select the VLAN ID to set its forwarding method.
- IPv6 Multicast Forwarding—Select one of the following options:
 - By MAC Address—Select to enable the MAC address method for forwarding Multicast packets.
 - By IPv6 Group Address—Select to enable the IPv6 group address method for forwarding Multicast packets.
 - By Source Specific IPv6 Group Address—Select to enable the source-specific IPv4 group address method for forwarding Multicast packets.
- 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.

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

IGMP/MLD Snooping

Multicast registration is the process of listening and responding to Multicast registration protocols. The available protocols are IGMP for IPv4 and MLD for IPv6.

When IGMP/MLD snooping is enabled in a device on a VLAN, it analyzes the IGMP/MLD 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/MLD messages to register to receive a Multicast stream, optionally from a specific source, the device adds the registration to the MFDB.

The following versions are supported:

- IGMP v1/v2/ v3
- MLD v1/v2

Note—The device supports IGMP/MLD Snooping only on static VLANs. It does not support IGMP/MLD Snooping on dynamic VLANs.

When IGMP/MLD Snooping is enabled globally or on a VLAN, all IGMP/MLD 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 (M routers) that are generating IGMP/MLD queries.
- Which ports are receiving PIM, DVMRP, or IGMP/MLD query protocols.

These VLANs are displayed on the IGMP/MLD Snooping page.

Ports, asking to join a specific Multicast group, issue an IGMP/MLD 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 support selective IPv6 Multicast forwarding, bridge Multicast filtering must be enabled (in the *Multicast > Feature Configuration* page), and MLD Snooping must be enabled globally and for each relevant VLAN in the MLD Snooping pages.

IGMP Snooping

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

1. Click *Configuration > Multicast > IGMP Snooping*.

LINKSYS LGS528 28-Port Gigabit Managed Switch Logout Help
Firmware Version: 1.1.0.27

System Status Quick Start Configuration Maintenance Support

System Management

- System Information
- TCAM Resources
- Management Session Timeout
- Time
- SNMP
- Logs

Port Management

- VLAN Management
- Spanning Tree Management
- MAC Address Management

Multicast

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

IP Interface

- IP Network Operations
- Security
- Access Control List
- Quality of Service

IGMP Snooping

IGMP Snooping: Enable

Apply Cancel

IGMP Snooping Table

| VLAN ID | IGMP Snooping Status | Router IGMP Version | Auto Learn MRouter Ports | IGMP Querier Status | IGMP Querier Version | IGMP Querier IP Address | Immediate Leave |
|-------------------------|----------------------|---------------------|--------------------------|---------------------|----------------------|-------------------------|-----------------|
| <input type="radio"/> 1 | Disabled | v3 | Enabled | Disabled | v2 | 192.168.1.120 | Disabled |

Edit

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.
3. Select a VLAN, and click **Edit**.

Edit IGMP Snooping

Select Your VLAN

VLAN ID:

VLAN Settings

IGMP Snooping Status: Enable

Auto Learn MRouter Ports: Enable Immediate Leave: Enable

IGMP Querier: Enable

IGMP Querier Version: IGMPv2 IGMPv3

Querier Source IP Address: Auto User Defined

Enter the parameters:

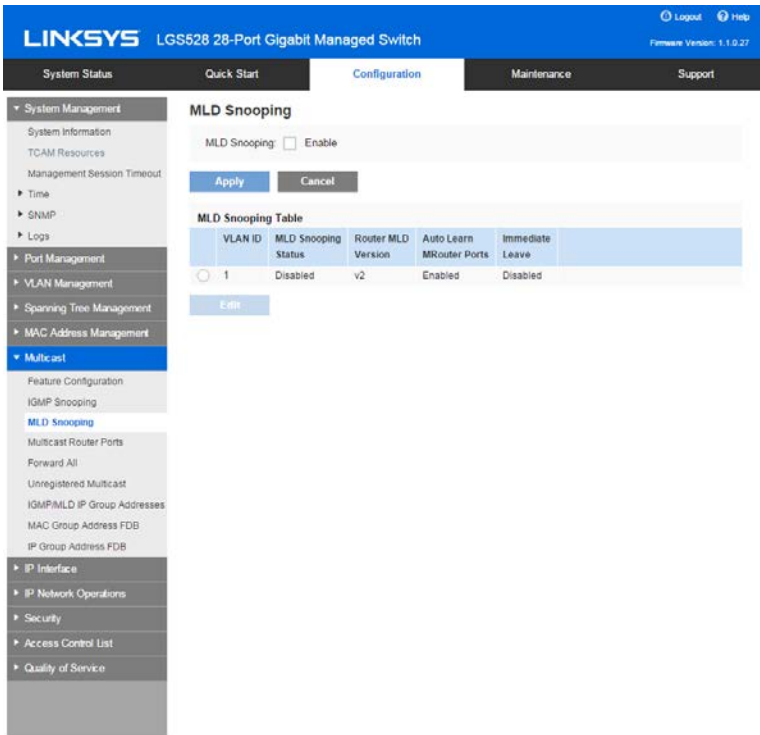
- VLAN ID-Select the VLAN ID on which IGMP snooping is defined.
- VLAN Settings
- IGMP Snooping Status-Enable or disable the monitoring of network traffic for the selected VLAN.
- Auto Learn MRouter Ports -Select to enable auto learning of the ports to which the Mrouter is connected.
- Immediate Leave-Select 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.

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

MLD Snooping

To enable MLD Snooping and configure it on a VLAN:

1. Click *Configuration > Multicast > MLD Snooping*.



When MLD Snooping is globally enabled, the device monitoring network traffic can determine which hosts have requested to receive Multicast traffic. The device performs MLD Snooping only if both MLD snooping and Bridge Multicast filtering are enabled.

2. Select MLD Snooping to enable the feature.

The following fields are displayed in the MLD Snooping Table:

- VLAN ID-VLAN ID on which MLD snooping is defined.
- MLD Snooping Status-Whether the monitoring of network traffic for the selected VLAN is enabled or disabled.
- Router MLD Version-MLD version supported.
- Auto Learn MRouter Ports -Whether auto learning of the ports to which the Mrouter is connected is enabled.

- Immediate Leave—Whether Immediate Leave to decrease the time it takes to block a Multicast stream sent to a member port when an MLD Group Leave message is received on that port is enabled.

3. Click **Edit**.

4. Enable or disable the following features for the selected VLAN:

- VLAN ID—Select a VLAN on which to configure MLD Snooping.
- MLD Snooping Status—Select to enable MLD snooping globally on all interfaces.
- Auto-Learn MRouter Ports—Select to enable Auto Learn of the Multicast router.
- Immediate Leave—Select to enable the switch to remove an interface that sends a leave message from the forwarding table without first sending out MAC-based general queries to the interface. When an MLD Leave Group message is received from a host, the system removes the host port from the table entry. After it relays the MLD queries from the Multicast router, it deletes entries periodically if it does not receive any MLD membership reports from the Multicast clients. When enabled, this feature reduces the time it takes to block unnecessary MLD traffic sent to a device port.

5. 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/MLD 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:

1. Click *Configuration > Multicast > Multicast Router Ports*.

2. Enter some or all of following query filter criteria:
 - VLAN ID—Select the VLAN ID for the router ports that are described.
 - IP Version—Select whether the IP group address is an IPv4 or IPv6 Address.
 - Interface Type—Select whether to display ports or LAGs.
3. Click **Search**. The interfaces matching the query criteria are displayed.
4. For each port or LAG, select its association type.

The options:

- 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 IGM/MLDP query. To enable the dynamic learning of Multicast router ports, go to the IGMP/MLD Snooping page.
 - Forbidden—This port is not to be configured as a Multicast router port, even if IGMP/MLD 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.
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.

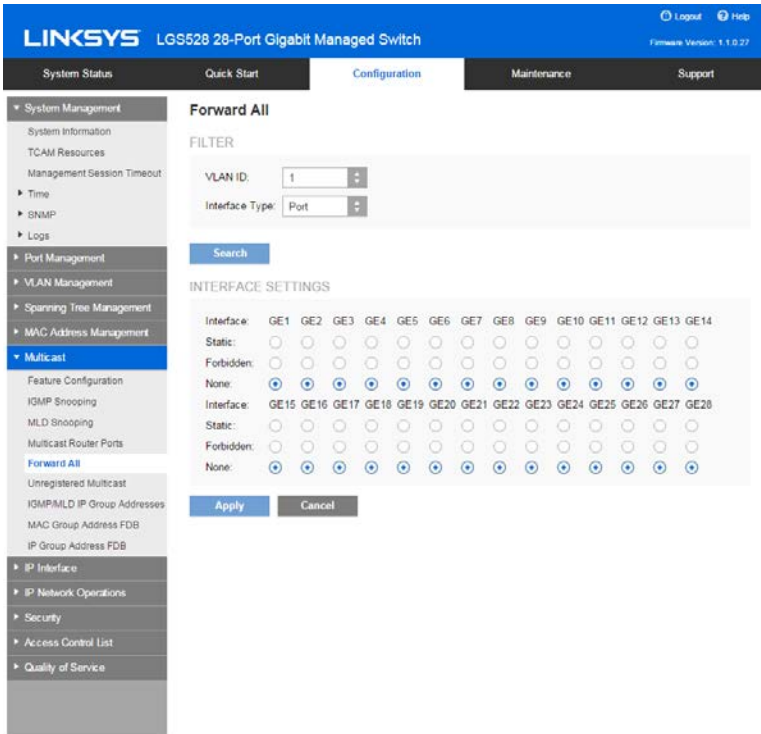
You can statically (manually) configure a port to Forward All, if the devices connecting to the port do not support IGMP/MLD.

IGMP/MLD 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:

1. Click *Configuration > Multicast > Forward All*.



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.
3. Click **Search**. The status of all ports/LAGs are displayed.
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/MLD snooping designated the port to join a Multicast group.
 - None—The port is not currently a Forward All port.
5. Click **Apply**. The Running Configuration file is updated.

Unregistered Multicast

Multicast frames are generally forwarded to all ports in the VLAN. If IGMP/MLD 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.

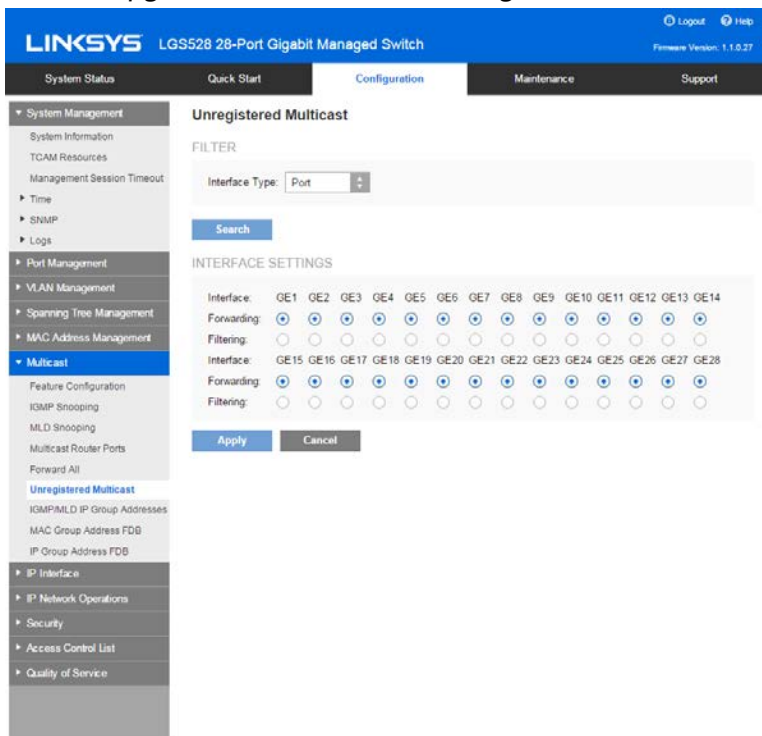
The Unregistered Multicast page enables handling Multicast frames that belong to groups that are not known to the device (unregistered Multicast groups). Unregistered Multicast frames are usually forwarded to all ports on the VLAN.

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:

1. Click *Configuration > Multicast > Unregistered Multicast*.



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

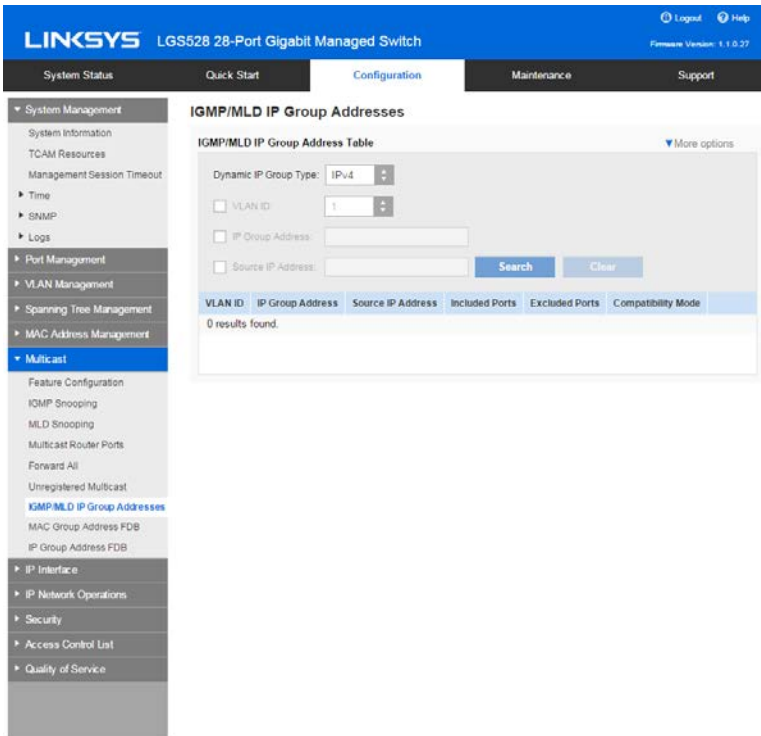
IGMP/MLD IP Group Addresses

The IGMP/MLD IP Group Addresses page displays the IPv4 group address learned from IGMP/MLD 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:

1. Click *Configuration > Multicast > IGMP/MLD IP Group Addresses*.



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

- Dynamic IP Group Type—Select whether IP group address is an IPv4 or IPv6 address.
- 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.

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/MLD 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/MLD 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/MLD 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:

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

The screenshot shows the Linksys web interface for an LGS528 28-Port Gigabit Managed Switch. The top navigation bar includes 'System Status', 'Quick Start', 'Configuration', 'Maintenance', and 'Support'. The left sidebar lists various configuration categories, with 'Multicast' selected. The main content area is titled 'MAC Group Address FDB' and contains a search form with fields for 'VLAN ID' (with a '(1-4094)' range) and 'MAC Group Address'. Below the search form is a table with columns for 'VLAN ID' and 'MAC Group Address', currently showing '0 results found'. At the bottom of the table are buttons for 'Add', 'Membership', and 'Delete'.

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

Add MAC Group Address

Enter New Group Address

VLAN ID: (1-4094)

MAC Group Address:

Apply **Close**

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.
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.
7. Click **Search** to display the port or LAG membership.
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.
9. Click **Apply**. 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:

1. Click *Configuration > Multicast > IP Group Address FDB*.

The screenshot shows the Linksys web interface for an LGS528 28-Port Gigabit Managed Switch. The top navigation bar includes 'System Status', 'Quick Start', 'Configuration', 'Maintenance', and 'Support'. The left sidebar lists various configuration categories, with 'Multicast' selected. The main content area is titled 'IP Group Address FDB' and contains a search form with fields for 'VLAN ID', 'IP Version', 'IP Group Address', and 'Source IP Address'. Below the search form is a table with columns for 'VLAN ID', 'IP Group Address', and 'Source IP Address', and a message indicating '0 results found'. At the bottom of the table are buttons for 'Add', 'Membership', and 'Delete'.

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

2. Enter the parameters required for filtering.
 - VLAN ID—Enter the VLAN ID of the group to be displayed.
 - IP Version—Select whether the IP group address is an IPv4 or IPv6 address.
 - 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.
3. Click **Search**. The results are displayed in the lower block.
 4. Click **Add** to add a static IP Multicast group address.
 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.
 6. Click **Apply**. The IP Multicast group is added, and the device is updated.
 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.
8. For each interface, select its association type.

The options:

 - 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.
 9. Click **Apply**. The Running Configuration file is updated.

Chapter 10 - IP Interface

This section describes IP interfaces and covers the following topics:

- [IPv4](#)
- [IPv6](#)

IPv4

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

- [Overview](#)
- [IPv4 Interface in Layer 2 System Mode](#)
- [IPv4 Interface in Layer 3 System Mode](#)
- [IPv4 Static Routes](#)
- [ARP](#)

Overview

Some features are only available in Layer 2 or Layer 3 system mode:

- 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:

- 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 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:

1. Click *Configuration > IP Interface > IPv4 > IPv4 Interfaces*.

The screenshot shows the Linksys LGS528 28-Port Gigabit Managed Switch web interface. The top navigation bar includes 'System Status', 'Quick Start', 'Configuration', 'Maintenance', and 'Support'. The left sidebar lists various management categories, with 'IP Interface' selected. The main content area is titled 'IPv4 Interface' and contains the following configuration fields:

- Management VLAN: 1
- IP Address Type: Dynamic (DHCP) Static IP Address
- Dynamic IP Address: Renew Now
- IP Address: 192.168.1.120
- IP Subnet Mask: Subnet Mask 255.255.255.0 Prefix Length
- User Defined Default Gateway: 192.168.1.1 Default Gateway: 192.168.1.1

Buttons for 'Apply' and 'Cancel' are located at the bottom of the configuration area.

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.*

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

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.

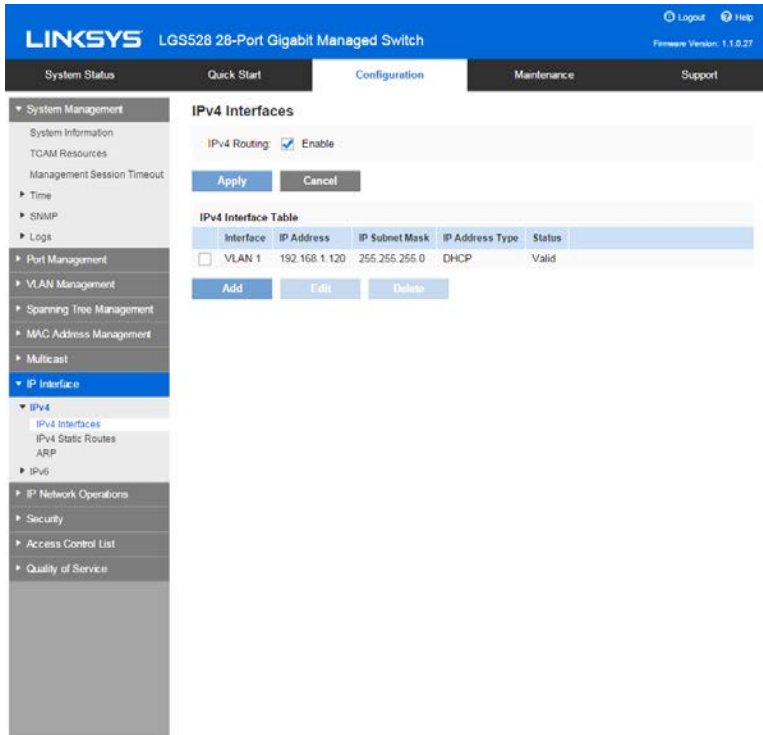
The IP address can be configured on a port, a LAG, or VLAN.

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:

1. Click *Configuration > IP Interface > IPv4 > IPv4 Interfaces*.



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

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

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:

1. Click *Configuration > Multicast > IPv4 > IPv4 Static Routes*.
2. Click **Add**.

Add IPv4 Static Route

Enter New Static Route

IP Subnet:

IP Subnet Mask: Subnet Mask

Prefix Length

Static Route Settings

Route Type: Reject Remote

Next Hop Router IP Address: Metric: (1-255)

Apply
Close

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.

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

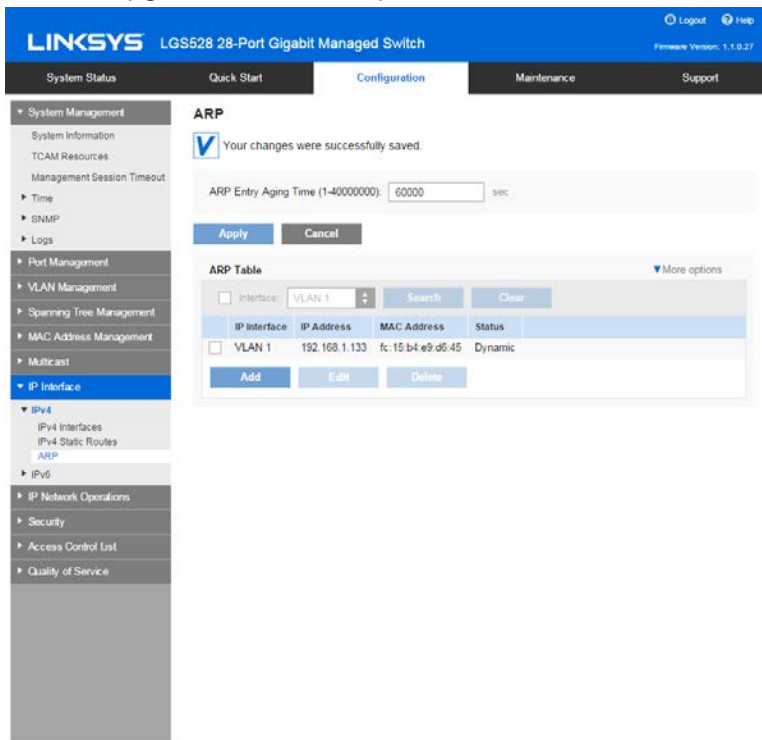
ARP

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:

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



The screenshot shows the configuration page for the ARP table on a Linksys LGS528 28-Port Gigabit Managed Switch. The page title is "LINKSYS LGS528 28-Port Gigabit Managed Switch" and the firmware version is 1.1.0.27. The navigation menu includes System Status, Quick Start, Configuration, Maintenance, and Support. The left sidebar shows the configuration tree with "IP Interface" selected. The main content area is titled "ARP" and displays a success message: "Your changes were successfully saved." Below this, there is a field for "ARP Entry Aging Time (1-40000000):" with a value of "60000" and a "sec" unit. There are "Apply" and "Cancel" buttons. The "ARP Table" section shows a table with columns for "Interface", "IP Address", "MAC Address", and "Status". The table contains one entry for "VLAN 1" with IP address "192.168.1.133" and MAC address "fc:15:b4:e9:d6:45", with a status of "Dynamic". There are "Add", "Edit", and "Delete" buttons below the table.

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.
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.

4. Click **Add**.

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.
5. 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 Interface](#)
- [IPv6 ISATAP Tunnel](#)

- [IPv6 Interface Addresses](#)
- [IPv6 Default Routers](#)
- [IPv6 Routes](#)
- [IPv6 Neighbors](#)

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 FE80:0000:0000:0000:9C00:876A: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 Static Routing

In the same way as occurs in IPv4 routing, frames addressed to the device's MAC address, but to an IPv6 address that is not known to the device, are forwarded to a next-hop device. This device may be the target end-station, or a router nearer the destination. The forwarding mechanism entails re-building a L2 frame around the (essentially) unchanged L3 packet received, with the next-hop device's MAC address as the destination MAC address.

The system uses Static Routing and Neighbor Discovery messages (similar to IPv4 ARP messages) to build the appropriate forwarding tables and next-hop addresses.

A route defines the path between two network devices. Routing entries added by the user are static, which are kept and used by the system until explicitly removed by the user, and are not changed by routing protocols. When static routes must be updated, this must be done explicitly by the user. It is the user's responsibility to prevent routing loops in the network.

Static IPv6 routes are either:

- Directly-attached, meaning that the destination is directly-attached to an interface on the device, so that the packet destination (which is the interface) is used as the next-hop address.

- Recursive, where only the next-hop is specified, and the outgoing interface is derived from the next-hop.

In the same manner, the MAC address of the next-hop devices (including directly-attached end-systems) are automatically derived using Network Discovery. However, the user may override and supplement this by manually adding entries to the Neighbors table.

IPv6 Interface

An IPv6 interface can be configured on a port, LAG, VLAN or tunnel.

As opposed to other types of interfaces, a tunnel interface is first created in the IPv6 Tunnel page, and then IPv6 interface is configured on the tunnel in this page.

To define an IPv6 interface:

1. Click *Configuration > IPv6 Interface > IPv6 > IPv6 Interface*.

| IPv6 Interface | Number of DAD Attempts | IPv6 Address Auto Configuration | Send ICMPv6 Messages |
|--------------------------------|------------------------|---------------------------------|----------------------|
| <input type="checkbox"/> VLAN1 | 1 | Enabled | Enabled |

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

Add IPv6 Interface

Enter New Interface

IPv6 Interface: Port LAG
 VLAN ISATAP Tunnel

Interface Settings

Number of DAD Attempts: (0-600)

IPv6 Address Auto Configuration: Enable

Send ICMPv6 Messages: Enable

Apply Close

3. Enter the fields:

- IPv6 Interface—Select a specific port, LAG, ISATAP tunnel or VLAN for the IPv6 address.
- 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 auto configuration from a DHCPv6 server.*

- Send ICMPv6 Messages—Enable generating unreachable destination messages.

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

IPv6 ISATAP Tunnel

Tunnels enable transmission of IPv6 packets over IPv4 networks. Each tunnel has a source IPv4 address and, if it is a manual tunnel, it also has a destination IPv4 address. The IPv6 packet is encapsulated between these addresses.

Note—Only the IPv6 management interface can be tunneled. To create an IPv6 tunnel, define an IPv6 interface as a tunnel in the IPv6 Interfaces page and continue configuring the tunnel in the IPv6 tunnel page.

ISATAP Tunnels

The type of tunnel that can be configured on the device is called an Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) tunnel, which is a point-to-multi-point tunnel. The source address is the IPv4 address (or one of the IPv4 addresses) of the device.

When configuring an ISATAP tunnel, the destination IPv4 address is provided by the router. Note the following:

- An IPv6 link local address is assigned to the ISATAP interface. The initial IP address is assigned to the interface, which is then activated.
- If an ISATAP interface is active, the ISATAP router IPv4 address is resolved via DNS by using ISATAP-to-IPv4 mapping. If the ISATAP DNS record is not resolved, the ISATAP host name-to-address mapping is searched for in the host mapping table.
- When the ISATAP router IPv4 address is not resolved via the DNS process, the ISATAP IP interface remains active. The system does not have a default router for ISATAP traffic until the DNS process is resolved.

Note—After configuring a tunnel, configure IPv6 interface in the IPv6 Interfaces page.

To configure an IPv6 tunnel:

1. Click *Configuration > IP Interface > IPv6 > IPv6 ISATAP Tunnel*.

The screenshot shows the configuration page for an IPv6 ISATAP Tunnel on a Linksys LGS528 28-Port Gigabit Managed Switch. The page title is "LINKSYS LGS528 28-Port Gigabit Managed Switch" with a firmware version of 1.1.0.27. The navigation menu includes System Status, Quick Start, Configuration, Maintenance, and Support. The left sidebar shows a tree view of configuration options, with "IP Interface" expanded to show "IPv6 ISATAP Tunnel" selected. The main content area is titled "IPv6 ISATAP Tunnel" and contains the following fields:

- Tunnel Number: 1
- Tunnel Type: ISATAP
- Source IPv4 Address: Auto None
- IPv4 Address: Interface: VLAN1
- ISATAP Router Name: Use Default User Defined:

At the bottom of the form are "Apply" and "Cancel" buttons.

2. Enter values for the following fields:

- Tunnel Number—Displays the automatic tunnel router domain number.
- Tunnel Type—Always ISATAP.
- Source IPv4 Address—The IPv4 address of the selected interface on the current device used to form part of the IPv6 address.
 - Auto—Automatically selects the lowest IPv4 address from among all of its configured IPv4 interfaces on the device. This option is equivalent to the Interface option in Layer 3, because in Layer 2 there is only one interface.

Note—If the IPv4 address is changed, the local address of the tunnel interface is also changed.

- None—Disable the tunnel.

- IPv4 Address—Enter the IPv4 source address to be used. <300-500>The IPv4 address configured must be one of the IPv4 addresses of the devices IPv4 interfaces.
- Interface—(In Layer 3) Select the IPv4 interface to be used.
- ISATAP Router Name—A global string that represents a specific automatic tunnel router domain name. The name can either be the default name (ISATAP) or a user-defined name.

Note—*The ISATAP tunnel is not operational if the underlying IPv4 interface is not in operation.*

3. Click **Apply**. The tunnel is saved to the Running Configuration file.

IPv6 Interface Addresses

To assign an IPv6 address to an IPv6 Interface:

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

The screenshot shows the Linksys web interface for an LGS528 28-Port Gigabit Managed Switch. The navigation menu on the left is expanded to 'IP Interface' > 'IPv6' > 'IPv6 Interface Addresses'. The main content area displays the 'IPv6 Interface Address Table' for interface GE1. The table contains one entry: a Link Local address (fe80::b575:eff:fe7c:5b7a) with a prefix length of 64, DAD Status of Tentative, and Type of System. There are 'Add' and 'Delete' buttons below the table.

| IPv6 Address Type | IPv6 Address | Prefix Length | DAD Status | Type |
|-------------------|--------------------------|---------------|------------|--------|
| Link Local | fe80::b575:eff:fe7c:5b7a | 64 | Tentative | System |

- To filter the table, select an interface name, and click **Search**. The interface appears in the IPv6 Address Table.
- Click **Add**.

The screenshot shows a configuration window titled "Add IPv6 Address". It contains the following fields and options:

- Enter New Address** (Section Header)
- IPv6 Interface:** A dropdown menu with "GE1" selected.
- IPv6 Address Type:** Two radio buttons: "Link Local Address" (selected) and "Global Address".
- IPv6 Address:** An empty text input field.
- Prefix Length:** An empty text input field with "(3-128)" as a hint.
- EUI-64:** A checkbox labeled "Enable" which is currently unchecked.
- At the bottom, there are two buttons: "Apply" (highlighted in blue) and "Close".

- 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 single IPv6 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.
- 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:

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

| Default Router IPv6 Address | IPv6 Interface | State | Type |
|--|----------------|-------|------|
| <input type="checkbox"/> fe80::4af9:b3ff:feca:de3f | VLAN1 | | |

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.

2. Click **Add** to add a static default router.

3. Enter the following fields:

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

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

IPv6 Routes

The IPv6 Forwarding Table contains the various routes that have been configured. One of these routes is a default route (IPv6 address:0) that uses the default router selected from the IPv6 Default Router List to send packets to destination devices that are not in the same IPv6 subnet as the device. In addition to the default route, the table also contains dynamic routes that are ICMP redirect routes received from IPv6 routers by using ICMP redirect messages. This could happen when the default router the device uses is not the router for traffic to which the IPv6 subnets that the device wants to communicate.

To view IPv6 routes:

1. Click *Configuration > IP Interface > IPv6 > IPv6 Routes*.

The screenshot shows the Linksys web interface for the LGS528 28-Port Gigabit Managed Switch. The navigation menu on the left includes System Management, Port Management, IP Interface, IP Network Operations, Security, Access Control List, and Quality of Service. The IP Interface menu is expanded to show IPv4 and IPv6 options. The IPv6 menu is further expanded to show IPv6 Interface, IPv6 ISATAP Tunnel, IPv6 Interface Addresses, IPv6 Default Routers, IPv6 Routes, and IPv6 Neighbors. The IPv6 Routes option is selected, displaying the IPv6 Routing Table.

| IPv6 Subnet Address | Prefix Length | IPv6 Interface | Next Hop Router IPv6 Address | Life Time | Route Type |
|----------------------|---------------|----------------|------------------------------|-----------|------------|
| :: | 0 | VLAN1 | fe80::4a6:b3ff:eca:de3f | | ND |
| 2001:540:c001:46c9:: | 64 | VLAN1 | :: | | ND |

This page displays the following fields:

- IPv6 Subnet Address—The IPv6 subnet address.
- Prefix Length—IP route prefix length for the destination IPv6 subnet address. It is preceded by a forward slash.
- IPv6 Interface—Interface used to forward the packet.
- Next Hop Router IPv6 Address—Address where the packet is forwarded.

Typically, this is the address of a neighboring router. It can be one of the following types.

- Link Local—An IPv6 interface and 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.

- Point-to-Point—A Point-to-point tunnel.
- Metric—Value used for comparing this route to other routes with the same destination in the IPv6 router table. All default routes have the same value.
- Life Time—Time period during which the packet can be sent, and resent, before being deleted.
- Route Type—How the destination is attached, and the method used to obtain the entry. The following values are:
 - Local—A directly-connected network whose prefix is derived from a manually-configured device's IPv6 address.
 - Dynamic—The destination is an indirectly-attached (remote) IPv6 subnet address. The entry was obtained dynamically via the ND or ICMP protocol.

IPv6 Neighbors

The IPv6 Neighbors page enables configuring and viewing the list of IPv6 neighbors on the IPv6 interface. The IPv6 Neighbor Table (also known as IPv6 Neighbor Discovery Cache) displays the MAC addresses of the IPv6 neighbors that are in the same IPv6 subnet as the device. This is the IPv6 equivalent of the IPv4 ARP Table. When the device needs to communicate with its neighbors, the device uses the IPv6 Neighbor Table to determine the MAC addresses based on their IPv6 addresses.

This page displays the neighbors that were automatically detected. Each entry displays to which interface the neighbor is connected, the neighbor's IPv6 and MAC addresses, the entry type, and the state of the neighbor.

To define IPv6 neighbors:

1. Click *Configuration > IP Interface > IPv6 > IPv6 Neighbors*.

The screenshot shows the web interface for a LINKSYS LGS528 28-Port Gigabit Managed Switch. The navigation menu includes System Status, Quick Start, Configuration, Maintenance, and Support. The left sidebar shows a tree view under System Management, with IP Interface selected. The main content area displays the IPv6 Neighbors configuration page, which includes an IPv6 Neighbor Table with the following data:

| Interface | IPv6 Address | MAC Address | Type | State |
|-----------|--|-------------------|---------|-------|
| VLAN1 | 2601:540:c001:4bc9:4af8:b3ff:feca:de3f | 48:f8:b3:ca:de:3f | Dynamic | Stale |
| VLAN1 | fe80::4af8:b3ff:feca:de3f | 48:f8:b3:ca:de:3f | Dynamic | Stale |

The following fields are displayed for the neighboring interfaces:

- IPv6 Interface—Neighboring IPv6 interface type.
- IPv6 Address—IPv6 address of a neighbor.
- MAC Address—MAC address mapped to the specified IPv6 address.
- Type—Neighbor discovery cache information entry type.
- State—Specifies the IPv6 neighbor status. The values are:
 - Incomplete—Address resolution is working. The neighbor has not yet responded.
 - Reachable—Neighbor is known to be reachable.
 - Stale—Previously-known neighbor is unreachable. No action is taken to verify its reachability until traffic must be sent.
 - Delay—Previously-known neighbor is unreachable. The interface is in Delay state for a predefined Delay Time. If no reachability confirmation is received, the state changes to Probe.
 - Probe—Neighbor is no longer known to be reachable, and Unicast Neighbor Solicitation probes are being sent to verify the reachability.

Chapter 11 - IP Network Operations

This section covers the following topics:

- [Domain Name System](#)
- [DHCP](#)
- [IP Source Guard](#)
- [DHCP Snooping Binding Database](#)
- [ARP Inspection](#)
- [Interface Settings](#)

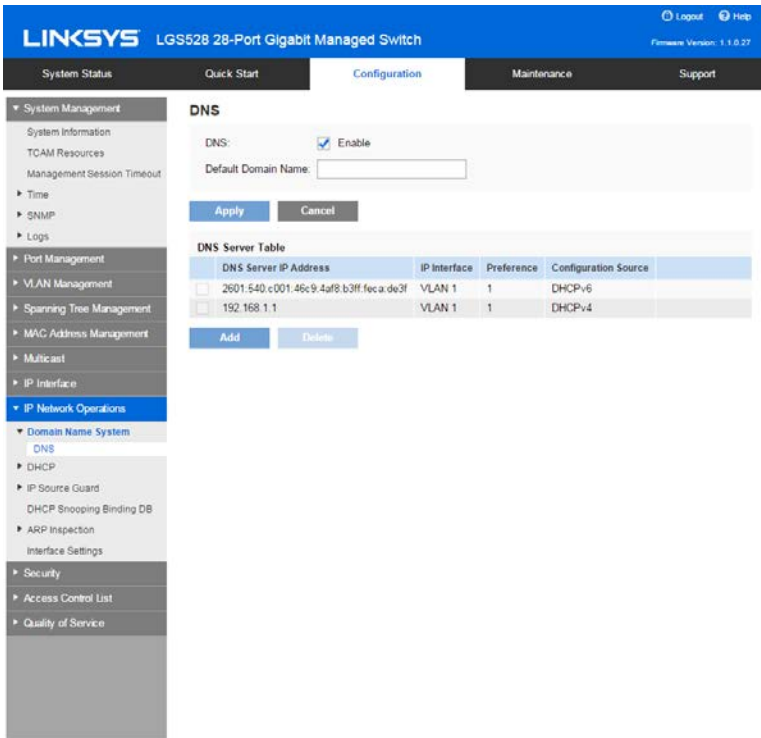
Domain Name System

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.

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

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



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.
- 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)

3. To add a DNS server, click **Add**. (Up to eight DNS servers can be defined.)

4. 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.

5. Click **Apply**. The DNS server is saved to the Running Configuration file.

DHCP

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.

Table 1: 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 |

Table 2: 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 |

Table 3: 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.



1. Device sends DHCPDISCOVER to request an IP address or DHCPREQUEST to accept an IP address and lease.
2. Device snoops packet and adds the IP-MAC information to the DHCP Snooping Binding database.
3. Device forwards DHCPDISCOVER or DHCPREQUEST packets.
4. DHCP server sends DHCP OFFER packet to offer an IP address, DHCPACK to assign one, or DHCPNAK to deny the address request.
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.
6. Device forwards DHCP OFFER, 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.

| 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 |
| DHCPRELEASE | Same as DHCPDECLINE. | Same as DHCPDECLINE. |
| DHCPINFORM | Forward to trusted interfaces only. | Forward to trusted interfaces only. |
| DHCPLEASEQUE RY | 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.

The following describes DHCP Snooping and DHCP Relay default options.

| Option | Default State |
|---------------------------------------|---------------|
| DHCP Snooping | Enabled |
| Option 82 Insertion | Disabled |
| Option 82 Passthrough | Disabled |
| Verify MAC Address | Enabled |
| Backup DHCP Snooping Binding Database | Disabled |
| DHCP Relay | Disabled |

Configuring DHCP Work Flow

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

1. Enable DHCP Snooping and/or DHCP Relay in the *Configuration > IP Network Operations > DHCP Relay & Snooping* page.
2. Define the interfaces on which DHCP Snooping is enabled in the *Configuration > IP Network Operations > DHCP > DHCP Interfaces* page.
3. Configure interfaces as trusted or untrusted in the *Configuration > IP Network Operations > DHCP > Trust Interface* page.
4. Optional. Add entries to the DHCP Snooping Binding database in the *Configuration > IP Network Operations > DHCP Snooping Binding Database* page.

DHCP Relay & Snooping

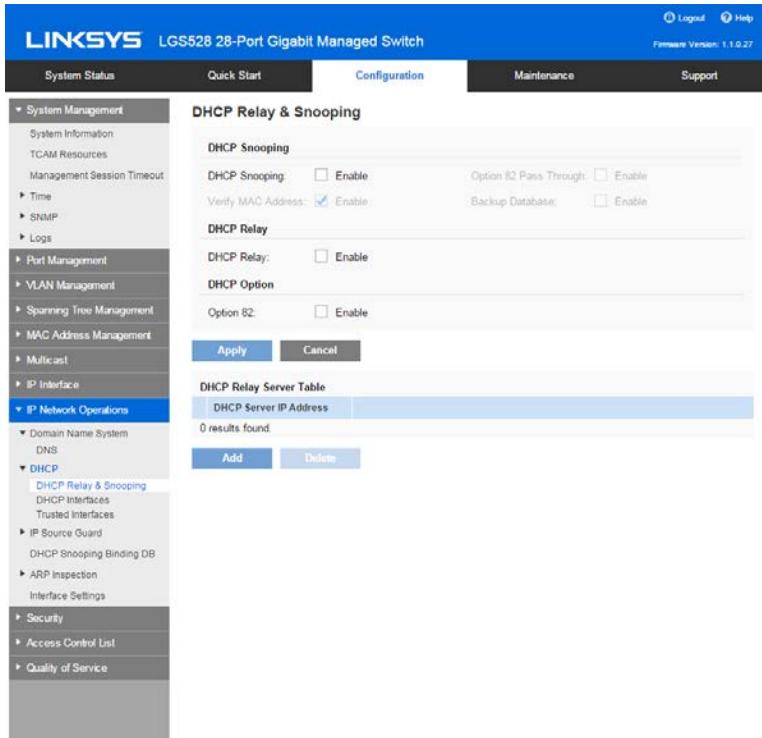
This section describes how the DHCP Relay and Snooping features are implemented via the Web-based interface.

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 Relay & Snooping:

1. Click *Configuration > IP Network Operations > DHCP > DHCP Relay & Snooping*.



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

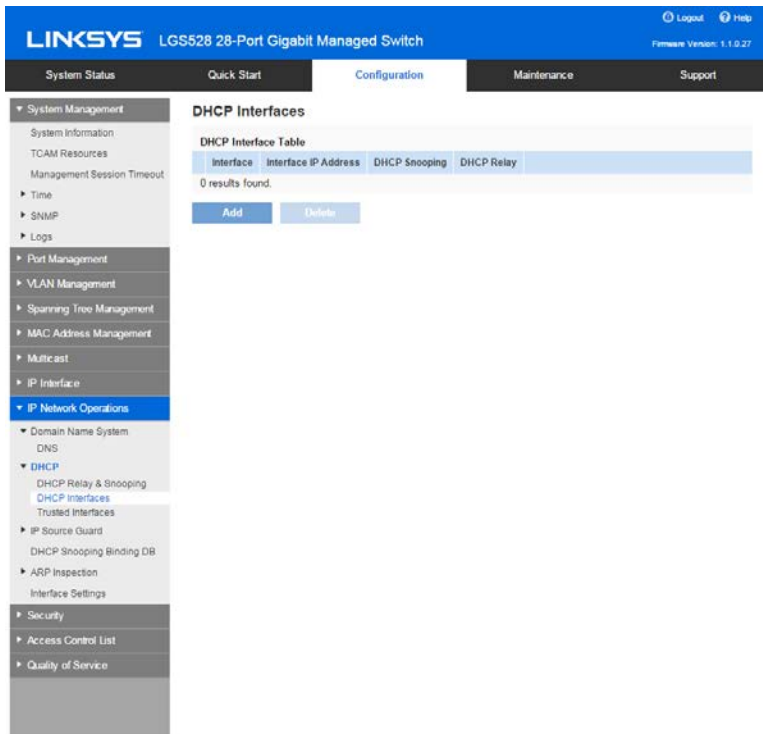
- DHCP Snooping—Select to enable DHCP Snooping.
- Option 82 Passthrough—Select to leave foreign Option 82 information when forwarding packets.
- Verify MAC Address—Select to verify that the source MAC address of the Layer 2 header matches the client hardware address as appears in the DHCP Header (part 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 Relay and Snooping on specific interfaces:

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

3. Select the interface and the features to be enabled: DHCP Relay or DHCP Snooping.
4. Click **Apply**. The settings are written to the Running Configuration file.

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 trusted.

To designate an interface as untrusted go to ARP Inspection ([p. 208](#)).

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:

1. Enable DHCP Snooping in the *IP Network Operations > DHCP > DHCP Relay & Snooping* page.

LINKSYS
LGS528 28-Port Gigabit Managed Switch
Logout Help
Firmware Version: 1.1.0.27

System Status
Quick Start
Configuration
Maintenance
Support

System Management

- System Information
- TCAM Resources
- Management Session Timeout
- Time
- SNMP
- Logs
- Port Management
- VLAN Management
- Spanning Tree Management
- MAC Address Management
- Multicast
- IP Interface
- IP Network Operations
- Domain Name System
 - DNS
- DHCP
 - DHCP Relay & Snooping
 - DHCP Interfaces
 - Trusted Interfaces
- IP Source Guard
 - DHCP Snooping Binding DB
- ARP Inspection
- Interface Settings
- Security
- Access Control List
- Quality of Service

DHCP Relay & Snooping

DHCP Snooping

DHCP Snooping: Enable Option 82 Pass Through: Enable

Verify MAC Address: Enable Backup Database: Enable

DHCP Relay

DHCP Relay: Enable

DHCP Option

Option 82: Enable

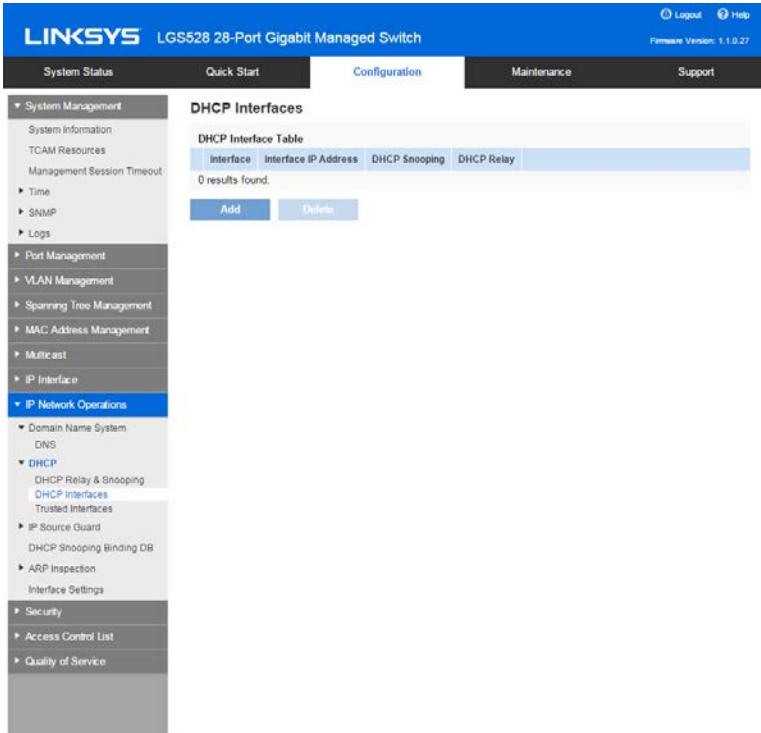
Apply
Cancel

DHCP Relay Server Table

| DHCP Server IP Address |
|------------------------|
| 0 results found |

Add
Delete

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



3. Configure interfaces as trusted or untrusted in the *IP Network Operations > DHCP > Trusted Interfaces* page.

LINKSYS LGS528 28-Port Gigabit Managed Switch Logout Help Firmware Version: 1.1.0.27

System Status **Quick Start** Configuration Maintenance Support

- System Management
 - System Information
 - TCAM Resources
 - Management Session Timeout
 - Time
 - SNMP
 - Logs
 - Port Management
 - VLAN Management
 - Spanning Tree Management
 - MAC Address Management
 - Multicast
 - IP Interface
 - IP Network Operations**
 - Domain Name System
 - DNS
 - DHCP
 - DHCP Relay & Snooping
 - DHCP Interfaces
 - Trusted Interfaces
 - IP Source Guard
 - Feature Configuration
 - IP Source Guard Interfaces
 - DHCP Snooping Binding DB
 - ARP Inspection
 - Interface Settings
 - Security
 - Access Control List
 - Quality of Service

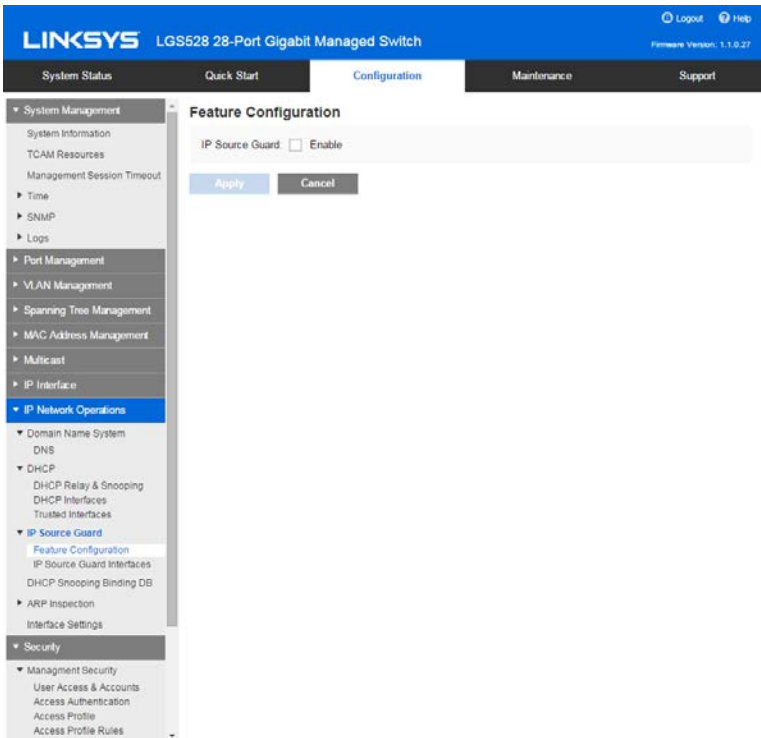
Interface Settings

Interface Settings Table

Interface Type: Port Search

| Interface | DHCP Snooping Trusted Interface | IP Source Guard | ARP Inspection Trusted Interface |
|----------------------------|---------------------------------|-----------------|----------------------------------|
| <input type="radio"/> GE1 | No | No | No |
| <input type="radio"/> GE2 | No | No | No |
| <input type="radio"/> GE3 | No | No | No |
| <input type="radio"/> GE4 | No | No | No |
| <input type="radio"/> GE5 | No | No | No |
| <input type="radio"/> GE6 | No | No | No |
| <input type="radio"/> GE7 | No | No | No |
| <input type="radio"/> GE8 | No | No | No |
| <input type="radio"/> GE9 | No | No | No |
| <input type="radio"/> GE10 | No | No | No |
| <input type="radio"/> GE11 | No | No | No |
| <input type="radio"/> GE12 | No | No | No |
| <input type="radio"/> GE13 | No | No | No |
| <input type="radio"/> GE14 | No | No | No |
| <input type="radio"/> GE15 | No | No | No |
| <input type="radio"/> GE16 | No | No | No |
| <input type="radio"/> GE17 | No | No | No |
| <input type="radio"/> GE18 | No | No | No |
| <input type="radio"/> GE19 | No | No | No |
| <input type="radio"/> GE20 | No | No | No |
| <input type="radio"/> GE21 | No | No | No |
| <input type="radio"/> GE22 | No | No | No |
| <input type="radio"/> GE23 | No | No | No |
| <input type="radio"/> GE24 | No | No | No |
| <input type="radio"/> GE25 | No | No | No |
| <input type="radio"/> GE26 | No | No | No |
| <input type="radio"/> GE27 | No | No | No |

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



5. Enable IP Source Guard on the untrusted interfaces as required in the *IP Network Operations > IP Source Guard > IP Source Guard Interfaces* page.

LINKSYS LGS528 28-Port Gigabit Managed Switch Logout Help
Firmware Version: 1.1.0.27

System Status Quick Start **Configuration** Maintenance Support

- System Management
 - System Information
 - TCAM Resources
 - Management Session Timeout
 - Time
 - SNMP
 - Logs
 - Port Management
 - VLAN Management
 - Spanning Tree Management
 - MAC Address Management
 - Multicast
 - IP Interface
 - IP Network Operations**
 - Domain Name System
 - DNS
 - DHCP
 - IP Source Guard
 - Feature Configuration
 - IP Source Guard Interfaces**
 - DHCP Snooping Binding DB
 - ARP Inspection
 - Interface Settings
 - Security
 - Access Control List
 - Quality of Service

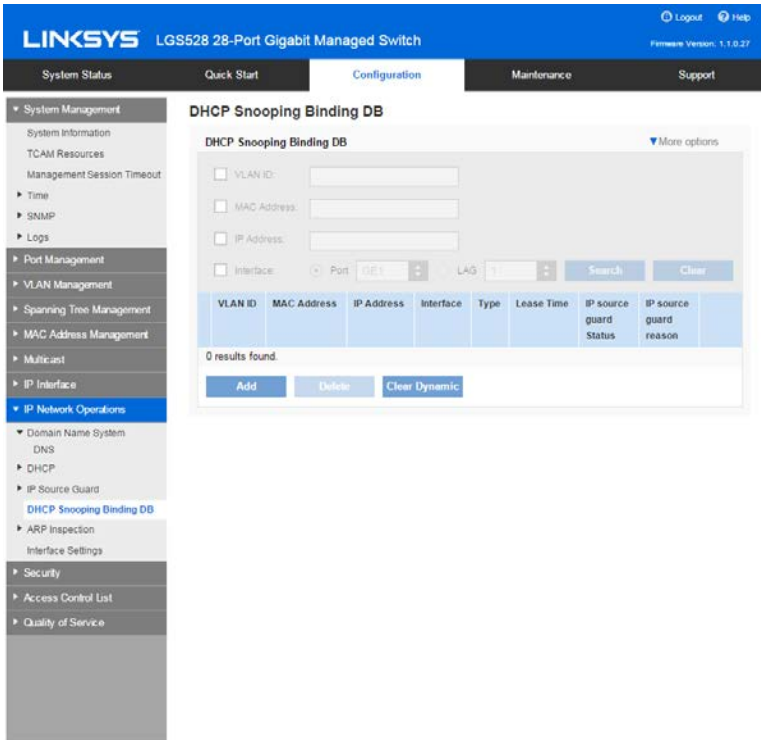
Interface Settings

Interface Settings Table

Interface Type:

| Interface | DHCP Snooping Trusted Interface | IP Source Guard | ARP Inspection Trusted Interface |
|----------------------------|---------------------------------|-----------------|----------------------------------|
| <input type="radio"/> GE1 | No | No | No |
| <input type="radio"/> GE2 | No | No | No |
| <input type="radio"/> GE3 | No | No | No |
| <input type="radio"/> GE4 | No | No | No |
| <input type="radio"/> GE5 | No | No | No |
| <input type="radio"/> GE6 | No | No | No |
| <input type="radio"/> GE7 | No | No | No |
| <input type="radio"/> GE8 | No | No | No |
| <input type="radio"/> GE9 | No | No | No |
| <input type="radio"/> GE10 | No | No | No |
| <input type="radio"/> GE11 | No | No | No |
| <input type="radio"/> GE12 | No | No | No |
| <input type="radio"/> GE13 | No | No | No |
| <input type="radio"/> GE14 | No | No | No |
| <input type="radio"/> GE15 | No | No | No |
| <input type="radio"/> GE16 | No | No | No |
| <input type="radio"/> GE17 | No | No | No |
| <input type="radio"/> GE18 | No | No | No |
| <input type="radio"/> GE19 | No | No | No |
| <input type="radio"/> GE20 | No | No | No |
| <input type="radio"/> GE21 | No | No | No |
| <input type="radio"/> GE22 | No | No | No |
| <input type="radio"/> GE23 | No | No | No |
| <input type="radio"/> GE24 | No | No | No |
| <input type="radio"/> GE25 | No | No | No |
| <input type="radio"/> GE26 | No | No | No |
| <input type="radio"/> GE27 | No | No | No |

6. View entries to the binding database in the *IP Network Operations > DHCP Snooping Binding BD* page.



Enabling IP Source Guard

To enable IP Source Guard globally:

Click *IP Network Operations > IP Source Guard > Feature Configuration*.

Select **Enable** to enable IP Source Guard globally.

Click **Apply** to enable IP Source Guard.

IP Source Guard Interfaces

If IP Source Guard is enabled on an untrusted port/LAG, 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* (p. 197) for more information about enabling IP Source Guard on interfaces.

To configure IP Source Guard on interfaces:

1. Click *IP Network Operations > IP Source Guard > IP Source Guard Interfaces*.
2. Select Port/LAG from the *Interface Type* field and click Search. The ports/LAGs on this unit are displayed along with the following:
 - IP Source Guard—Indicates whether IP Source Guard is enabled on the port.
 - DHCP Snooping Trusted Interface—Indicates whether this is a DHCP trusted interface.
3. Select the port/LAG and click **Edit**.

Edit Interface Settings

Select Your Interface

Interface: Port GE13 LAG 1

Interface Settings

DHCP Snooping Trusted Interface: Yes No

ARP Inspection Trusted Interface: Yes No

IP Source Guard: Enable

Apply **Close**

4. Select *Enable* in the *IP Source Guard* field.
5. Click **Apply** to copy the setting to the Running Configuration file.

DHCP Snooping Binding Database

View the DHCP Snooping 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 [DHCPv4 Relay & Snooping \(p. 194\)](#).

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:

1. Click *IP Network Operations > DHCP Snooping Binding DB*.

The entries in the Binding database are displayed:

- VLAN ID—VLAN on which packet is expected.
- MAC Address—MAC address to be matched.
- IPv4 Address—IP address to be matched.
- Interface—Interface on which packet is expected.
- Type—Displays whether entry is dynamic or static.
- Lease Time—If the entry is dynamic, displays the amount of time that the entry is to be active in the database. If there is no Lease Time, it is infinite.
- IP Source Guard Status—Displays whether the interface is active.
- IP Source Guard Reason—If the interface is not active, displays the reason.

The following reasons are possible:

- No Problem—Interface is active.
- No Snoop VLAN—DHCP Snooping is not enabled on the VLAN.
- Trusted Port—Port has become trusted.
- Resource Problem—TCAM resources are exhausted.

2. To see a subset of these entries, enter the relevant search criteria and click **Search**. To add an entry, click **Add**.

Add to DHCP Snooping Binding Database

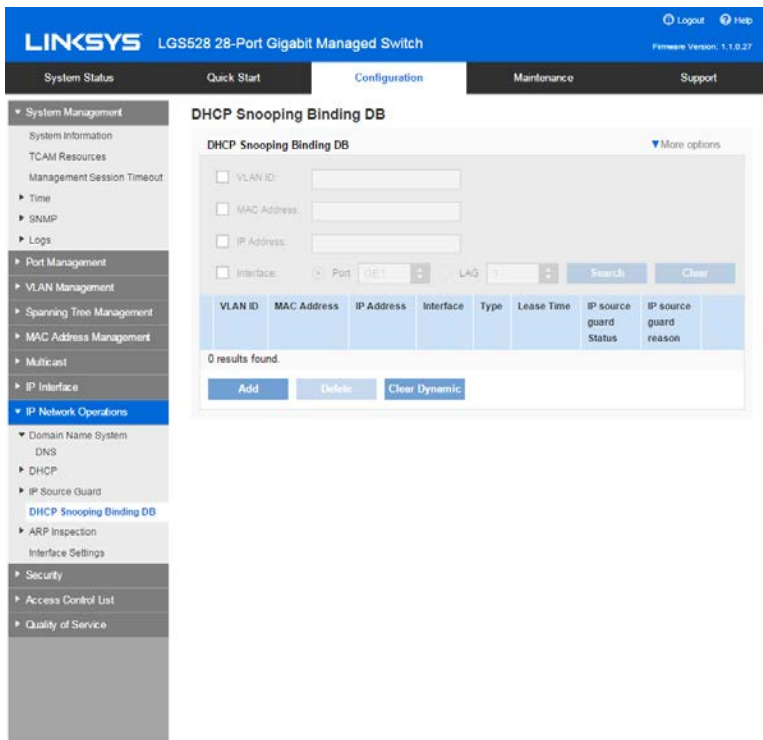
See How the DHCP Snooping Binding Database is Built (p. 192) for a description of how dynamic entries are added to the DHCP Snooping Binding database.

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. Note that if the IP source guard and/or ARP inspection features are active, the clients that are not written in the DHCP Snooping Binding database are not be able to connect to the network.

To add entries to the DHCP Snooping Binding database:

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 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

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.)
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 a 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.

How ARP Prevents Cache Poisoning

The ARP inspection feature relates to interfaces as either trusted or untrusted.

See *IP Network Operations > ARP Inspection > ARP Inspection Interface* 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 ARP 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 (*ARP Inspection > Feature Configuration* 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.

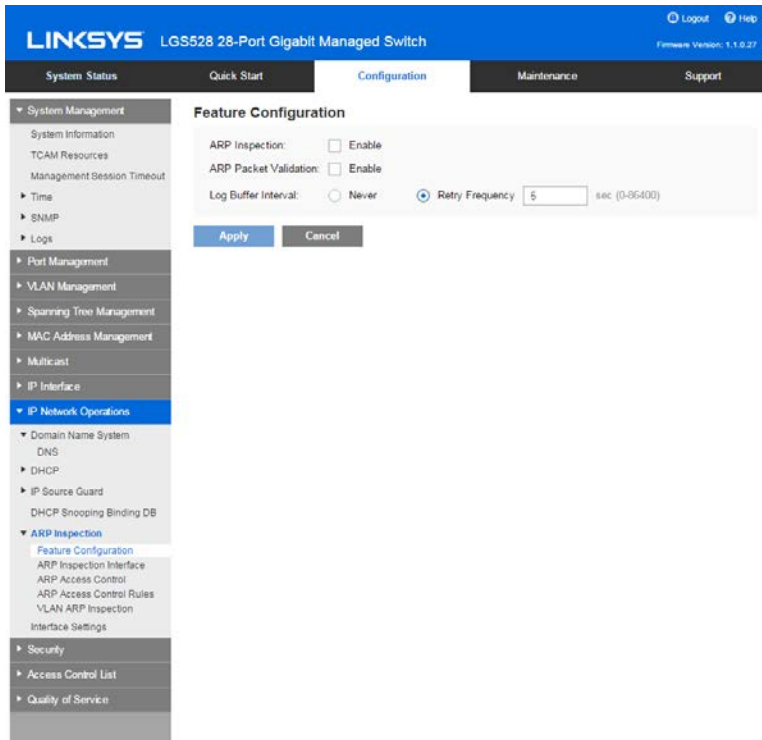
ARP Defaults

| Option | Default State |
|--------------------------------|--|
| Dynamic ARP Inspection | Disabled |
| ARP Packet Validation | Disabled |
| ARP Inspection Enabled on VLAN | Disabled |
| Log Buffer Interval | SYSLOG message generation for dropped packets is enabled at 5 seconds interval |

Feature Configuration

To configure ARP Inspection:

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



2. Configure interfaces as ARP trusted or untrusted in the *Configuration > IP Network Operations > ARP Inspection > ARP Inspection Interface* page.
3. Add rules in the *Configuration > IP Network Operations > ARP Inspection > ARP Access Control Rules* pages.
4. Define the VLANs on which ARP Inspection is enabled and the Access Control Rules for each VLAN in the *Configuration > IP Network Operations > ARP Inspection > VLAN ARP Inspection* page.

ARP Inspection Interface

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 DB* page).

By default, ports/LAGs are ARP Inspection untrusted.

To change the ARP trusted status of a port/LAG:

1. Click *Configuration > IP Network Operations > ARP Inspection > ARP Inspection Interface*.
2. To set a port/LAG as trusted or untrusted, select the port/LAG and click **Edit**.
3. Select *Yes* or *No* for *DHCP Snooping Trusted Interface*.
4. Select *Yes* or *No* for *ARP Inspection Trusted Interface*.
5. Click **Apply** to save the settings to the Running Configuration file.

To define ARP Inspection Properties

1. Click *Configuration > IP Network Operations > ARP Inspection > Feature Configuration*.
2. Enter the following fields:
 - ARP Inspection—Select to enable ARP Inspection.
 - ARP Packet Validation—Select to enable the following validation checks:
 - Source MAC — Compares the packets source MAC address in the Ethernet header against the senders MAC address in the ARP request. This check is performed on both ARP requests and responses.
 - Destination MAC — Compares the packets destination MAC address in the Ethernet header against the destination interfaces 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.
 - Log Buffer Interval—Select one of the following options:
 - Retry Frequency—Enable sending SYSLOG messages for dropped packets. Enter the frequency with which the messages are sent.
 - Never—Disabled SYSLOG dropped packet messages.
3. Click **Apply**. The settings are defined, and the Running Configuration file is updated.

ARP Access Control

To add a new access control list with a single rule to the ARP Inspection table:

1. Click *Configuration > IP Network Operations > ARP Inspection > ARP Access Control*.
2. To add a rule, click **Add**.

Add ARP Access Control

Enter New Access Control

ARP Access Control Name:

Enter New Rule

IP Address:

MAC Address:

3. Enter the fields:
 - *ARP Access Control Name*—Enter a user-created name.
 - *IP Address*—IP address of packet.
 - *MAC Address*—MAC address of packet.
4. Click **Apply**. The settings are defined, and the Running Configuration file is updated.

ARP Access Control Rules

To add more rules to a previously created ARP Access Control group:

1. Click Configuration > IP Network Operations > ARP Inspection > ARP Access Control Rules.
2. Click **Add**.
3. Select an Access Control Group and enter the fields:
 - *IP Address*—IP address of packet.
 - *MAC Address*—MAC address of packet.
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:

1. Click *Configuration > IP Network Operation > ARP Inspection > VLAN ARP Inspection*.
2. Move the VLAN from the *Available VLANs* list to the *ARP Inspection Enabled* list.
3. To associate an ARP Access Control group with a VLAN, click **Add**.

The screenshot shows a configuration dialog box titled "Add VLAN ARP Access Control". It is divided into two sections: "Select Your VLAN" and "Enter New Access Control". In the "Select Your VLAN" section, there is a "VLAN ID:" label followed by a text input field and a small dropdown arrow icon. In the "Enter New Access Control" section, there is an "ARP Access Control Name:" label followed by a text input field and a small dropdown arrow icon. At the bottom of the dialog, there are two buttons: "Apply" (in blue) and "Close" (in grey).

4. Select the VLAN number and select a previously defined ARP Access Control Name group.
5. Click **Apply**. The settings are defined, and the Running Configuration file is updated.

Interface Settings

To configure trusted interfaces:

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

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

This chapter covers the following sections:

- [Management Security](#)
- [RADIUS](#)
- [Network Access Control](#)
- [Port Security](#)
- [Storm Control](#)

Management Security

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.

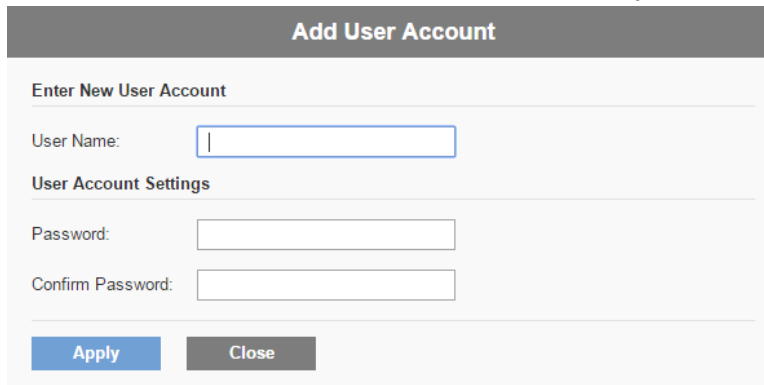
If an authentication method fails or the user has insufficient privilege level, the user is denied access to the device. In other words, if authentication fails at an authentication method, the device stops the authentication attempt; it does not continue and does not attempt to use the next authentication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

Notes

- *After adding a user (as described below), the default user is removed from the system.*
- *It is not permitted to delete all users. If all users are selected, the Delete button is disabled.*

To add a new user:

1. Click *Configuration > Security > Management Security > User Access & Accounts*.
2. 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.
3. Click **Add** to add a new user or click **Edit** to modify a user.



The screenshot shows a web interface for adding a user account. The title is "Add User Account". Below the title is a subtitle "Enter New User Account". There are three input fields: "User Name", "Password", and "Confirm Password". At the bottom, there are two buttons: "Apply" and "Close".

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

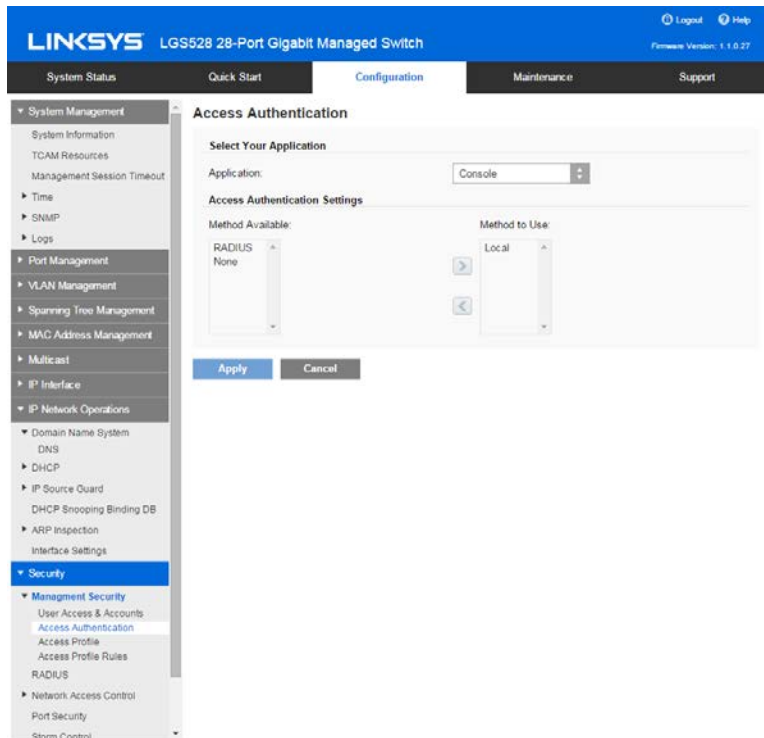
Access Authentication

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

To define authentication methods for an access method:

1. Click *Configuration > Security > Management Security > Access Authentication*.



2. Select an access method from the *Application* list.
3. Use the arrows to move the authentication method between the *Method Available* column and the *Method To Use* column. The first method selected is the first method that is used.
 - RADIUS—User is authenticated on a RADIUS server. You must have configured one or more RADIUS servers.
 - None—User is allowed to access the device without authentication.
 - Local—Username and password are checked against the data stored on the local device. These username and password pairs are defined in the User Accounts page.

Note—The *Local* or *None* authentication method must always be selected last. All authentication methods selected after *Local* or *None* are ignored.

4. Click **Apply**. The selected authentication methods are associated with the access method.

Access Profile

Access profiles determine how to authenticate and authorize users accessing the device through various access methods. Access Profiles can limit management access from specific sources.

Only users who pass both the active access profile and are authorized based on the authentication methods that correspond to the Access Method are given management access to the device. For more information, see Management Access Authentication.

There can only be a single access profile active on the device at one time. Access profiles consist of one or more rules. The rules are executed in order of their priority within the access profile (top to bottom).

Rules are composed of filters that include the following elements:

- **Access Methods**—Methods for accessing and managing the device: The authentication method for the selected access method is specified in Management Access Authentication
 - Telnet
 - Hypertext Transfer Protocol (HTTP)
 - Secure HTTP (HTTPS)
 - Simple Network Management Protocol (SNMP)
 - All of the above
- **Action**—Permit or deny access to an interface or source address.
- **Interface**—Which ports, LAGs, or VLANs are permitted to access or are denied access to the web-based configuration utility.
- **Source IP Address**—IP addresses or subnets. Access to management methods might differ among user groups. For example, one user group might be able to access the device module only by using an HTTPS session, while another user group might be able to access the device module by using both HTTPS and Telnet sessions.

The Access Profile page displays the access profiles that are defined and enables selecting one access profile to be the active one.

When a user attempts to access the device through an access method, the device looks to see if the active access profile explicitly permits management access to the device through this method. If no match is found, access is denied.

When an attempt to access the device is in violation of the active access profile, the device generates a SYSLOG message to alert the system administrator of the attempt.

If a console-only access profile has been activated, the only way to deactivate it is through a direct connection from the management station to the physical console port on the device.

For more information, see [Access Profile Rules \(p. 220\)](#).

Use the Access Profiles page to create an access profile and to add its first rule. If the access profile only contains a single rule, you are finished. To add additional rules to the profile, use the Profile Rules page.

1. Click *Configuration > Security > Management Security Management > Access Profile*.
2. To change the active access profile, select a profile from the *Active Access Profile* drop-down menu and click **Apply**. This makes the chosen profile the active access profile.

Note—A caution message appears if you selected *Console Only*. If you continue, you are immediately disconnected from the web-based configuration utility and can access the device only through the console port. This only applies to device types that offer a console port.

3. Click **OK** to select the active access profile or click **Cancel** to discontinue the action.
4. Click **Add** to open the *Add Access Profile* page. The page allows you to configure a new profile and one rule.

The screenshot shows the 'Add Access Profile' configuration page. It is titled 'Add Access Profile' and is divided into two sections: 'Enter New Profile' and 'Enter New Rule'. The 'Enter New Profile' section has a text input field for 'Access Profile Name'. The 'Enter New Rule' section contains several configuration options: 'Rule Priority' (input field with '(1-65535)' next to it), 'Management Access Method' (radio buttons for All, Telnet, HTTP, HTTPS, and SNMP), 'Access Control' (radio buttons for Permit and Deny), 'Interface' (radio buttons for All, Port, LAG, and VLAN), 'Source IP Address' (radio buttons for All and User Defined), 'IP Version' (radio buttons for IPv4 and IPv6), 'IP Address' (input field), 'IP Subnet Mask' (radio buttons for Subnet Mask and Prefix Length), and corresponding input fields for the selected options. At the bottom of the form are 'Apply' and 'Close' buttons.

5. Enter the Access Profile Name. This name can contain up to 32 characters.
6. Enter the new rule parameters.
 - Rule Priority—Enter the rule priority. When the packet is matched to a rule, user groups are either granted or denied access to the device. The rule priority is essential to matching packets to rules, as packets are matched on a first-match basis. One is the highest priority.
 - Management Access Method—Select the management method for which the rule is defined. The options are:
 - All—Assigns all management methods to the rule.
 - Telnet—Users requesting access to the device that meets the Telnet access profile criteria are permitted or denied access.
 - HTTP—Users requesting access to the device that meets the HTTP access profile criteria, are permitted or denied.
 - Secure HTTP (HTTPS)—Users requesting access to the device that meets the HTTPS access profile criteria, are permitted or denied.
 - SNMP—Users requesting access to the device that meets the SNMP access profile criteria are permitted or denied.
 - Access Control—Select the action attached to the rule. The options are:
 - Permit—Permits access to the device if the user matches the settings in the profile.
 - Deny—Denies access to the device if the user matches the settings in the profile.
 - Interface—Select the interface attached to the rule. The options are:
 - All—Applies to all ports, VLANs, and LAGs.
 - Port—Rule applies to ports.
 - LAG—Rule applies to LAGs.
 - VLAN—Rule applies to VLANs.
 - Source IP Address—Select the type of source IP address to which the access profile applies. The Source IP Address field is valid for a subnetwork. Select one of the following values:
 - All—Applies to all types of IP addresses.
 - User Defined—Applies to only those types of IP addresses defined in the fields.
 - IP Version—Enter the version of the source IP address: Version 6 or Version 4.

- IP Address—Enter the source IP address.
 - IP Subnet Mask—Select the format for the subnet mask for the source IP address, and enter a value in one of the fields:
 - Network Mask—Select the subnet to which the source IP address belongs and enter the subnet mask in dotted decimal format.
 - Prefix Length—Select the Prefix Length and enter the number of bits that comprise the source IP address prefix.
7. Click **Apply**. The access profile is written to the Running Configuration file. You can now select this access profile as the active access profile.

Access Profile Rules

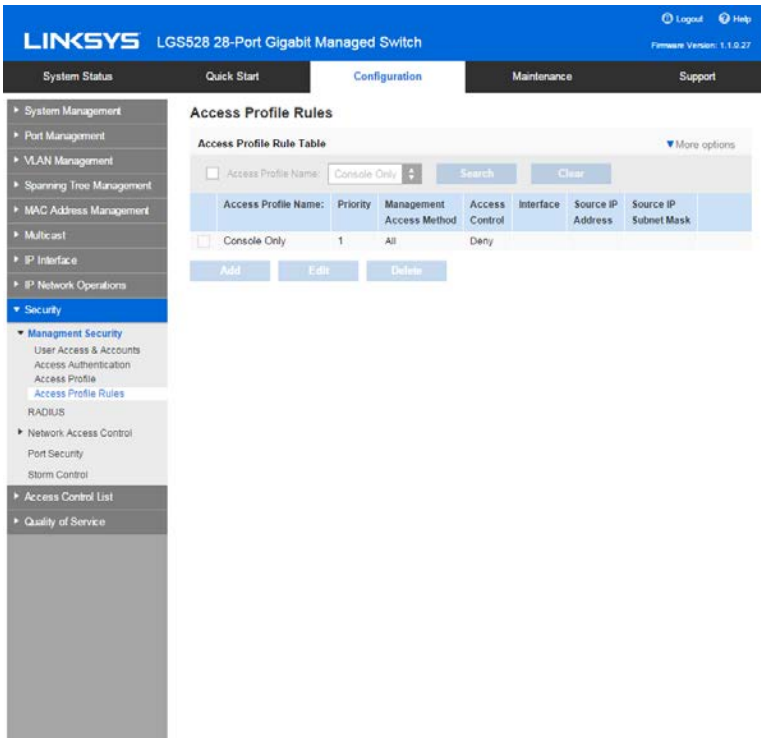
Access profiles can contain up to 128 rules to determine who is permitted to manage and access the device, and the access methods that may be used.

Each rule in an access profile contains an action and criteria (one or more parameters) to match. Each rule has a priority; rules with the lowest priority are checked first. If the incoming packet matches a rule, the action associated with the rule is performed. If no matching rule is found within the active access profile, the packet is dropped.

For example, you can limit access to the device from all IP addresses except IP addresses that are allocated to the IT management center. In this way, the device can still be managed and has gained another layer of security.

To add profile rules to an access profile:

1. Click *Configuration > Security > Management Security > Access Profile Rules*.



2. Select the Filter field, and an access profile. Click **Search**.

The selected access profile appears in the *Profile Rule Table*.

3. Click **Add** to add a rule.

4. Enter the parameters.

- Access Profile Name—Select an access profile.
- Rule Priority—Enter the rule priority. When the packet is matched to a rule, user groups are either granted or denied access to the device. The rule priority is essential to matching packets to rules, as packets are matched on a first-fit basis.
- Management Access Method—Select the management method for which the rule is defined. The options are:
 - All—Assigns all management methods to the rule.
 - Telnet—Users requesting access to the device that meets the Telnet access profile criteria are permitted or denied access.
 - HTTP—Assigns HTTP access to the rule. Users requesting access to the device that meets the HTTP access profile criteria, are permitted or denied.
 - Secure HTTP (HTTPS)—Users requesting access to the device that meets the HTTPS access profile criteria, are permitted or denied.

- SNMP—Users requesting access to the device that meets the SNMP access profile criteria are permitted or denied.
- Access Control—Select Permit to permit the users that attempt to access the device by using the configured access method from the interface and IP source defined in this rule. Or select Deny to deny access.
- Interface—Select the interface attached to the rule. The options are:
 - All—Applies to all ports, VLANs, and LAGs.
 - Port—Select the port attached to the rule.
 - LAG—Select the LAG attached to the rule.
 - VLAN—Select the VLAN attached to the rule.
- Source IP Address—Select the type of source IP address to which the access profile applies. The Source IP Address field is valid for a subnetwork. Select one of the following values:
 - All—Applies to all types of IP addresses.
 - User Defined—Applies to only those types of IP addresses defined in the fields.
- IP Version—Select the supported IP version of the source address: IPv6 or IPv4.
- IP Address—Enter the source IP address.
- IP Subnet Mask—Select the format for the subnet mask for the source IP address, and enter a value in one of the field:
 - Network Mask—Select the subnet to which the source IP address belongs and enter the subnet mask in dotted decimal format.
 - Prefix Length—Select the Prefix Length and enter the number of bits that comprise the source IP address prefix.

5. Click **Apply**, and the rule is added to the access profile.

RADIUS

Remote Authorization Dial-In User Service (RADIUS) servers provide a centralized 802 .1X 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 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 by using 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.
- Accounting—Enable accounting of login sessions using the RADIUS server. This enables a system administrator to generate accounting reports from the RADIUS server.

Accounting Using a RADIUS Server

The user can enable accounting of login sessions using a RADIUS server.

The user-configurable, TCP port used for RADIUS server accounting is the same TCP port that is used for RADIUS server authentication and authorization.

The following defaults are relevant to this feature:

- No default RADIUS server is defined by default.
- If you configure a RADIUS server, the accounting feature is disabled by default.

Interactions with other features:

- None.

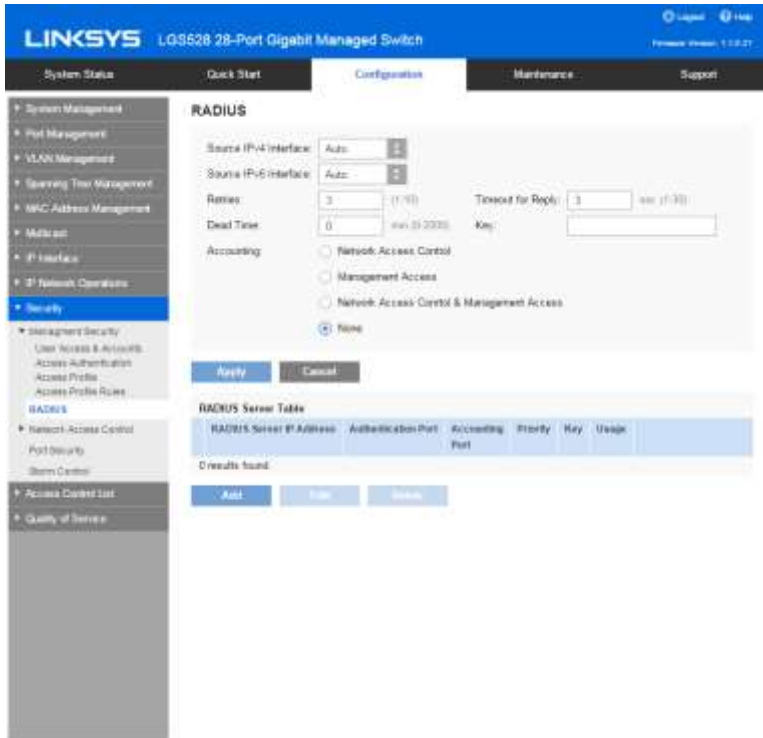
To user a RADIUS server, do the following:

1. Open an account for the device on the RADIUS server.
2. Configure that server along with the other parameters in the RADIUS and ADD RADIUS Server 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:

1. Click Configuration > Security > RADIUS.



2. Enter the default RADIUS parameters if required. If a value is not entered for a specific server (in the Add RADIUS Server page) the device uses the values in these fields.
 - Source IPv4 Interface—(In layer 3) Select the device IPv4 source interface to be used in messages for communication with the RADIUS server.
 - Source IPv6 Interface—(In layer 3) 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 String**—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.

- **RADIUS Accounting option**—Select one of the following options:
 - **Network Access Control**—Specifies that the RADIUS server is used for 802.1x port accounting.
 - **Management Access**—Specifies that the RADIUS server is used for user login accounting.
 - **Both Network Access Control and Management Access**—Specifies that the RADIUS server is used for both user login accounting and 802.1x port accounting.
 - **None**—Specifies that the RADIUS server is not used for accounting.

3. Click **Apply**. The RADIUS default settings for the device are updated in the Running Configuration file.

To add a RADIUS server, click **Add**.

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:

- 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 IP Name—Enter the name of the RADIUS server.
 - Authentication Port—Enter the UDP port number of the RADIUS server port for authentication requests.
 - Accounting Port—Enter the UDP port number of the RADIUS server port for accounting 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. If Use Default is selected, the device attempts to authenticate to the RADIUS server by using the default Key String.
 - Usage Type—Enter the RADIUS server authentication type. The options are:
 - 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.
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 publicly-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



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.

When a port is authorized, untagged and tagged traffic from all hosts connected to the port is bridged, based on the static VLAN membership port configuration.

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.



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.



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 unauthenticated VLAN.

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 Assignment

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.

| Authentication Method | Single-Host | Multi-Host | Multi-Sessions | |
|-----------------------|-------------|------------|----------------|--------------|
| | | | Device in L3 | Device in L2 |
| 802.1x | † | † | N/S | † |
| MAC | † | † | N/S | † |

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

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

To enable 802.1x authentication on a port:

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

The screenshot shows the configuration page for a Linksys LGS528 28-Port Gigabit Managed Switch. The page is titled "LINKSYS LGS528 28-Port Gigabit Managed Switch" and "Feature Configuration". The navigation menu on the left includes System Management, Port Management, VLAN Management, Spanning Tree Management, MAC Address Management, Multicast, IP Interface, IP Network Operations, Security, Management Security, Network Access Control, Access Control List, and Quality of Service. The Security section is expanded, and Network Access Control is selected. The Feature Configuration section has the following options:

- Port-Based Authentication: Enable
- Authentication Method: RADIUS, None, RADIUS, None
- Guest VLAN: Enable, Guest VLAN ID:

Buttons for "Apply" and "Cancel" are visible. Below the configuration options is a "VLAN Authentication Table" with columns for VLAN ID, VLAN Name, and Authentication. It shows "0 results found" and an "Edit" button.

2. Enable Port-based Authentication.
3. Select the Authentication Method.
4. Click **Apply**, and the Running Configuration file is updated.
5. Click *Configuration > Security > Network Access Control > Port Authentication*.

6. Select the required port and click **Edit**.

Edit Port Authentication

Select Your Interface

Interface: GE13

Interface Settings

Port Control: Force Unauthorized Auto Force Authorized

Host Authentication Mode: Multiple Host (802.1x) Multiple Sessions

RADIUS VLAN Assignment: Enable Guest VLAN: Enable

802.1x Authentication: Enable MAC Based Authentication: Enable

Periodic Reauthentication: Enable

Reauthentication Period: 3600 sec (300-4294967295)

Apply **Close**

7. Set the Host Authentication mode.
8. Set the *Administrative Port Control* field to *Auto*.
9. Define the authentication methods.
10. Click **Apply**, and the Running Configuration file is updated.

To configure 802.1x-based authentication:

1. Click *Configuration > Security > Network Access Control > Port Authentication*.
2. Select the required port and click **Edit**.
3. Enter the fields required for the port.

The fields in this page are described in [Port Authentication](#)

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

To configure the guest VLAN:

1. Click *Security > Network Access Control > Feature Configuration*.
2. Select Enable in the Guest VLAN field.
3. Select the guest VLAN in the Guest VLAN ID field.
4. Click **Apply**, and the Running Configuration file is updated.

To configure unauthenticated VLANs:

1. Click *Security > Network Access Control > Feature Configuration*.
2. Select a VLAN, and click **Edit**.
3. Select a VLAN.
4. Optionally, uncheck *Authentication* to make the VLAN an unauthenticated VLAN.
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:

1. Click *Configuration > Security > Network Access Control > Feature Configuration*.
2. Enter the parameters:
 - Port-Based Authentication—Enable or disable port-based authentication. If this is disabled 802 .1X 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.
 - 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.
3. Click **Apply**. The settings are written to the Running Configuration file.
4. To enable authentication for a specific VLAN, select it in the main page and click Edit.
5. Enable Authentication for that VLAN.
6. Click **Apply**. The settings are written to the Running Configuration file.

Port Authentication

The Port Authentication page enables configuration of 802.1X parameters for each 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 controls to their previous state.

Note—A port with 802.1x defined on it cannot become a member of a LAG.

To configure 802.1X authentication:

1. Click *Configuration > Security > Network Access Control > Port Authentication*.

This page displays authentication settings for all ports.

The *Current Port Control* displays 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*.

2. Select a port, and click **Edit**.

3. Enter the parameters:

- Interface—Select a port.
- Port Control—Select the *Administrative Port Authorization* state.

The options:

- Force Unauthorized—Denies the interface access by moving the interface into the unauthorized state. The device does not provide authentication services to the client through the interface.
- Auto—Enables port-based authentication and authorization on the device. The interface moves between an authorized or unauthorized state based on the authentication exchange between the device and the client.
- Force Authorized—Authorizes the interface without 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. The options are:
 - Selected—Enables using a guest VLAN for unauthorized ports. If a guest VLAN is enabled, the unauthorized port automatically joins the VLAN selected in the Guest VLAN ID field in the 802.1X Port Authentication page.

- After an authentication failure, and if guest VLAN is activated globally on a given port, the guest VLAN is automatically assigned to the unauthorized ports as an Untagged VLAN.
- Cleared—Disables guest VLAN on the port.
- 802.1X Based Authentication—Select to enable 802.1X authentication on the port.
- MAC Based Authentication—Select to enable MAC-based authentication on the port. The 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—Select to enable port re-authentication attempts after the specified Reauthentication Period.
- Reauthentication Period—Enter the number of seconds after which the selected port is reauthenticated.

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

Authenticated Hosts

To display details about authenticated users:

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 Mode and Port Mode Support

| Authentication Method | Multi-Host | Multi-Sessions | |
|-----------------------|------------|----------------|--------------|
| | | Device in L3 | Device in L2 |
| 802.1x | † | † | † |
| MAC | † | † | † |

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

Mode Behavior

| | Unauthenticated Traffic | | | |
|----------------------------|--|--|---|--|
| | With Guest VLAN | | Without Guest VLAN | |
| | Untagged | Tagged | Untagged | Tagged |
| Multi-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 |
| Lite Multi-sessions | N/S | N/S | Frames are dropped | Frames are dropped unless they belong to the unauthenticated VLANs |
| Full Multi-sessions | 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 |
| | | | | |
| | Authenticated Traffic | | | |
| | With RADIUS VLAN | | Without RADIUS VLAN | |
| | Untagged | Tagged | Untagged | Tagged |
| Multi-host | 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 |
| Lite Multi-sessions | 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 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 |

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:

1. Click *Configuration > Security > Port Security*.
2. Select an interface to be modified, and click **Edit**.
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.
 - Trap—Enable trap and set the trap frequency.
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:

1. Click *Configuration > Security > Storm Control*.
2. Select a port and click Edit.
3. Enter the parameters.
 - Port—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.
4. Click **Apply**. Storm control is modified, and the Running Configuration file is updated.

Chapter 13 - 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](#)
- [MAC-Based ACL/ACE](#)
- [IPv4-Based ACL/ACE](#)
- [IPv6-Based ACE/ACL](#)
- [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/IPv6- 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).

To create ACLs and associate them with an interface:

1. Create one or more of the following types of ACLs:
 - MAC-based ACL by using the *MAC Based ACL* page and the *MAC Based ACE* page.
 - IPv4-based ACL by using the *IPv4 Based ACL* page and the *IPv4 Based ACE* page.
 - IPv6-based ACL by using the *IPv6 Based ACL* page and the *IPv6 Based ACE* page.
2. Associate the ACL with interfaces by using the *ACL Binding* page.

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 page.

Only then can the ACL be modified, as described in this section.

MAC-Based ACL/ACE

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:

1. Click *Configuration > Access Control List > MAC Based ACL*.
This page contains a list of all currently-defined MAC-based ACLs.
2. Click **Add**.
3. Enter the name of the new ACL in the *ACL Name* field. ACL names are case-sensitive.
4. Click **Apply**. The MAC-based ACL is saved to the Running Configuration file.

To add rules (ACEs) to an ACL:

1. Click *Configuration > Access Control List > MAC Based ACE*.
2. Select an ACL, and click **Search**. The ACEs in the ACL are listed.
3. Click **Add**.

Add MAC Based ACE

Enter New ACE

ACL Name:

ACE Priority: (1 - 2147483647)

ACE Settings

Action on Matched Packets: Permit Deny Shutdown

Destination MAC Address: Any User Defined

Destination MAC Address Value:

Destination MAC Wildcard Mask: (0s for matching, 1s for non matching)

Source MAC Address: Any User Defined

Source MAC Address Value:

Source MAC Wildcard Mask: (0s for matching, 1s for non matching)

VLAN ID: (1-4094)

802.1p: Match

802.1p Value: (0-7) 802.1p Mask: (0-7)

EtherType: (5DD-FFFF)

Apply
Close

4. Enter the parameters.

- ACL Name—Select the name of the ACL to which an ACE is being added.
- 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.
5. Click **Apply**. The MAC-based ACE is saved to the Running Configuration file.

IPv4-Based ACL/ACE

IPv4-based ACLs are used to check IPv4 packets, while other types of frames, such as ARPs, are not checked.

Fields that 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.

To define an IPv4-based ACL:

1. Click *Configuration > Access Control List > IPv4 Based ACL*.
This page contains all currently defined IPv4-based ACLs.
2. Click **Add**.
3. Enter the name of the new ACL in the ACL Name field. The names are case-sensitive.
4. Click **Apply**. The IPv4-based ACL is saved to the Running Configuration file.

To add rules (ACEs) to an IPv4-based ACL:

1. Click *Configuration > Access Control List > IPv4-Based ACE*.
2. Select an ACL, and click **Search**. All currently-defined IP ACEs for the selected ACL are displayed.
3. Click **Add**.

The screenshot shows the 'Add IPv4 Based ACE' configuration page. The 'Enter New ACE' section contains the following fields and options:

- ACL Name:** ACL
- ACE Priority:** 1 (range 1 - 2147483647)
- ACE Settings:**
 - Action on Match Packets:** Permit, Deny, Shutdown
 - Protocol:** Any IPv4, Protocol List (ICMP), Protocol ID
 - Source IP Address:** Any, User Defined
 - Destination IP Address:** Any, User Defined
 - Source Port:** Any, Port (0-65535)
 - Destination Port:** Any, Port (0-65535)
 - Type of Services:** Any, DSCP (0-63), IP Precedence (0-7)

Buttons at the bottom: **Apply** and **Close**.

4. Enter the parameters.
 - **ACL Name**—Displays the name of the ACL.
 - **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 Services 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.
- Click **Apply**. The IPv4-based ACE is saved to the Running Configuration file.

IPv6-Based ACL/ACE

To define an IPv6-based ACL:

1. Click *Configuration > Access Control List > IPv6 Based ACL*.
This page contains all currently defined IPv6-based ACLs.
2. Click **Add**.
3. Enter the name of the new ACL in the ACL Name field. The names are case-sensitive.
4. Click **Apply**. The IPv6-based ACL is saved to the Running Configuration file.

To add rules (ACEs) to an IPv6-based ACL:

1. Click *Configuration > Access Control List > IPv6-Based ACE*.
2. Select an ACL, and click **Search**. All currently-defined IP ACEs for the selected ACL are displayed.

3. Click **Add**.

Add IPv6 Based ACE

Enter New ACE

ACL Name: IPv6 Based ACL
Priority: (1 - 2147483647)

ACE Settings

Action on Match: Permit Deny Shutdown

Protocol: Any (IPv6) Protocol List TCP Protocol ID

Source IP Address: Any User Defined

Source IP Address Value:
Source IP Prefix Length:

Destination IP Address: Any User Defined

Destination IP Address Value:
Destination IP Prefix Length:

Source Port: Any Port (0-65535)

Destination Port: Any Port (0-65535)

Type of Services: Any DSCP (0-63) IP Precedence (0-7)

4. Enter the parameters.

- **ACL Name**—Displays the name of the ACL.
- **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 IPv6 to accept all IP protocols. Otherwise select one of the following protocols from the drop-down list:
 - **ICMP**—Internet Control Message Protocol
 - **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 addresses 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 Prefix Length—Enter the prefix length of the source IP address.
- Destination IP Address—Select Any if all destination addresses 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 Prefix Length—Enter the prefix length of the destination IP address.
- 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 Services—The service type of the IP packet.
 - Any—Any service type
 - DSCP to Match—Differentiated Services Code Point (DSCP) to match
 - IP Precedence—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.

5. Click **Apply**. The IPv6-based ACE is saved to the Running Configuration file.

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:

1. Click *Configuration > Access Control List > ACL Binding*.
2. Select an interface type (Port or LAG).
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.*

4. Select an interface, and click **Edit**.
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.
 - IPv6 Based ACL—Select an IPv6-based ACL to be bound to the interface.
 - Permit Any Unmatched Packets—Select to enable/disable this action.
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

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](#)
- [CoS/802.1p to Queue](#)
- [DSCP to Queue](#)
- [Bandwidth Control](#)
- [Egress Shaping](#)
- [Basic QoS](#)
- [Advanced QoS](#)
- [QoS Statistics](#)

Overview

The QoS feature is used to optimize network performance. QoS classifies incoming traffic into traffic classes, based on attributes:

- Device configuration
- Ingress interface
- Packet content
- Combination of these attributes

QoS includes:

- Traffic Classification—Classifies each incoming packet as belonging to a specific traffic flow, based on the packet contents and/or the port. The classification is done by ACL (Access Control List), and only traffic that meets the ACL criteria is subject to CoS or QoS classification.

- **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

When using the QoS feature, all traffic of the same class receives the same treatment, which consists of a single QoS action of determining the egress queue on the egress port, based on the indicated QoS value in the incoming frame. This is 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 type of header field to be trusted is entered in the Basic QoS page. For every value of that field, an egress queue is assigned, indicating through which queue 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).

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 Basic QoS 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.

To configure general QoS parameters:

1. Choose the QoS mode (Basic, Advanced, or Disabled) by using the Feature Configuration page. The following steps in the workflow, assume that you have chosen to enable QoS.
2. Assign each interface a default CoS priority by using the Feature Configuration page.
3. Assign the schedule method (Strict Priority or WRR) and bandwidth allocation for WRR to the egress queues by using the Queue Scheduling page.
4. Set the trusted mode in the Basic QoS page or the advanced mode Feature Configuration page.
5. 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 IP packets are put into the egress queues based on their DSCP/TC value.
6. Designate an egress queue to each CoS/802.1p priority. If the device is in CoS/802.1p 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.
7. Enter bandwidth and rate limits in the following pages:
 - Set egress shaping per queue by using the *Egress Shaping* page.

- Set ingress rate limit and egress shaping rate per port by using the *Bandwidth Control* page.
8. Configure the selected mode:
- Configure Basic mode, as described in *To Configure Basic QoS Mode*
 - Configure Advanced mode, as described in *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:

1. Click *Configuration > Quality of Service > Feature Configuration*.

The screenshot shows the Linksys web interface for the LGS528 28-Port Gigabit Managed Switch. The top navigation bar includes 'System Status', 'Quick Start', 'Configuration', 'Maintenance', and 'Support'. The left sidebar menu is expanded to 'Quality of Service', with sub-items like 'Feature Configuration', 'Queue Scheduling', 'CoS802.1p to Queue', 'DSCP to Queue', 'Bandwidth Control', 'Egress Shaping', 'Basic QoS', 'Advanced QoS', and 'QoS Statistics'. The 'Feature Configuration' page displays three radio button options: 'Disable', 'Basic QoS' (which is selected), and 'Advanced QoS'. Below these are 'Apply' and 'Cancel' buttons. The 'Interface CoS Table' section has a dropdown menu set to 'Port' and a 'Search' button. The table below lists 28 interfaces (GE1 to GE28) with a 'Default CoS' column, all of which are currently set to 0. At the bottom of the table are 'Edit' and 'Refresh' buttons.

2. Set the QoS mode. The following options are available:
 - Disable—QoS is disabled on the device.

- Basic—QoS is enabled on the device in Basic mode.
 - Advanced—QoS is enabled on the device in Advanced mode.
3. Select Port/LAG and click **Search** to display/modify all ports/LAGs on the device and their CoS information.
 4. 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 *Basic QoS* page.
 5. Click **Apply**. The Running Configuration file is updated.

To set QoS on an interface:

1. Select an interface and click **Edit**.
2. 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).
3. Click **Apply**. The interface default CoS value is saved to Running Configuration file.

Queue Scheduling

The device supports four queues for each interface. Queue 4 is the highest priority queue. Queue 1 is the lowest priority queue.

There are two ways of determining how traffic in queues is handled:

- 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).

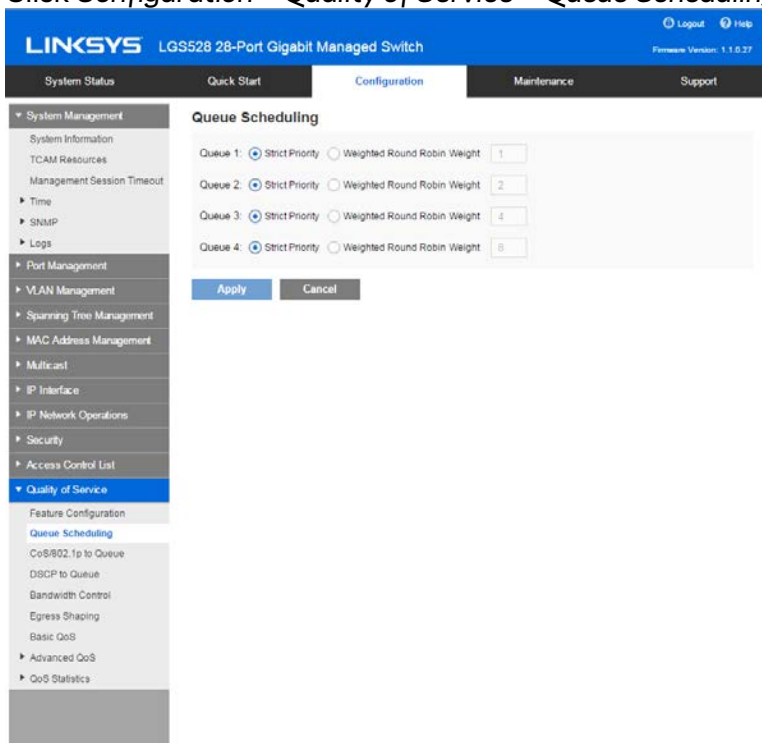
The queuing modes can be selected in the Queue Scheduling page. When the queuing mode is by strict priority, the priority sets the order in which queues are serviced, starting with Queue 4 (the highest priority queue) and going to the next lower queue when each queue is completed.

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:

1. Click *Configuration > Quality of Service > Queue Scheduling*.



2. Enter the parameters.
 - Queue—Displays the queue number.
 - Scheduling Method: Select one of the following options:
 - Strict Priority—Traffic scheduling for the selected queue and all higher queues is based strictly on the queue priority.

- WRR—Traffic scheduling for the selected queue is based on WRR. The period time is divided between the WRR queues that are not empty, meaning they have descriptors to egress. This happens only if strict priority queues are empty.
 - WRR Weight—If WRR is selected, enter the WRR weight assigned to the queue.
 - % of WRR Bandwidth—Displays the amount of bandwidth assigned to the queue. These values represent the percent of the WRR weight.
3. Click **Apply**. The queues are configured, and the Running Configuration file is updated.

CoS/802.1p to Queue

The CoS/802.1p to Queue page maps 802.1p priorities to egress queues. The CoS/802.1p to Queue Table determines the egress queues of the incoming packets based on the 802.1p priority in their VLAN Tags. For incoming untagged packets, the 802.1p priority is the default CoS/802.1p priority assigned to the ingress ports.

The following table describes the default mapping:

| 802.1p Values (0-7, 7 being the highest) | Queue (4 queues 1-4, 4 being the highest priority) | Notes |
|---|---|---------------------------------------|
| 0 | 1 | Background |
| 1 | 1 | Best Effort |
| 2 | 2 | Excellent Effort |
| 3 | 3 | Critical Application - VoIP phone SIP |
| 4 | 3 | Video |
| 5 | 4 | Voice - Cisco IP phone default |
| 6 | 4 | Interwork Control - VoIP phone RTP |
| 7 | 4 | Network Control |

By changing the CoS/802.1p to Queue mapping (CoS/802.1p to Queue page), the Queue schedule method (Queue Scheduling page) and bandwidth allocation (Bandwidth page), it is possible to achieve the desired quality of service in a network.

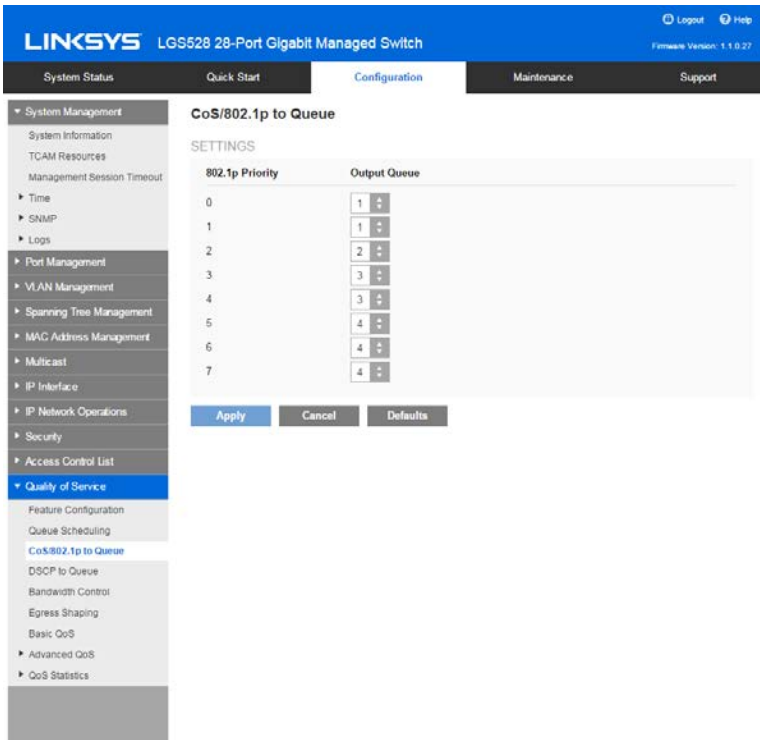
The CoS/802.1p to Queue mapping is applicable only if one of the following exists:

- The device is in QoS Basic mode and CoS/802.1p trusted mode
- The device is in QoS Advanced mode and the packets belong to flows that are CoS/802.1p trusted.

Queue 1 has the lowest priority; queue 4 has the highest priority.

To map CoS/802.1p values to egress queues:

1. Click Configuration>Quality of Service > General > CoS/802.1p to Queue.



2. Enter the parameters.

- 802.1p—Displays the 802.1p priority tag values to be assigned to an egress queue, where 0 is the lowest and 7 is the highest priority.
- Output Queue—Select the egress queue to which the 802.1p priority is mapped. Either four or eight egress queues are supported, where Queue 4 is the highest priority egress queue and Queue 1 is the lowest priority.

3. For each 802.1p priority, select the Output Queue to which it is mapped.

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

DSCP to Queue

The DSCP (IP Differentiated Services Code Point) to Queue page maps DSCP values to egress queues. The DSCP to Queue Table determines the egress queues of the incoming IP packets based on their DSCP values. The original VPT (VLAN Priority Tag) of the packet is unchanged. By simply changing the DSCP to Queue mapping and the Queue schedule method and bandwidth allocation, it is possible to achieve the desired quality of services in a network.

The DSCP to Queue mapping is applicable to IP packets if:

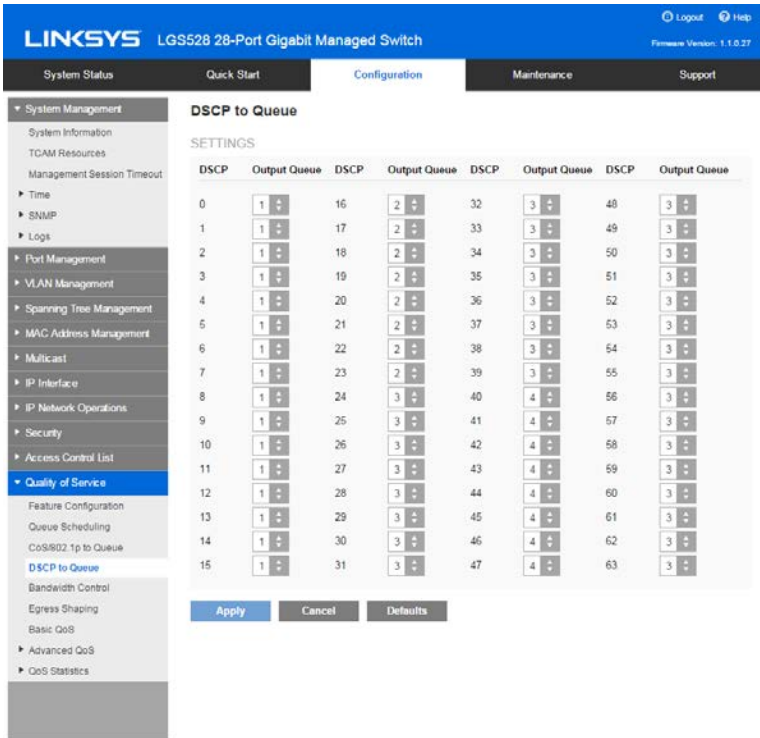
- The device is in QoS Basic mode and DSCP is the trusted mode, or
- The device is in QoS Advanced mode and packets belong to flows that are DSCP trusted

Non-IP packets are always classified to the best-effort queue.

| | | | | | | | | |
|-------|---------|---------|---------|----------|----------|----------|----------|-------|
| DSCP | 63 | 55 | 47 | 39 | 31 | 23 | 15 | 7 |
| Queue | 3 | 3 | 4 | 3 | 3 | 2 | 1 | 1 |
| DSCP | 62 | 54 | 46(EF) | 38(AF3) | 30(AF33) | 22(AF23) | 14 | 6 |
| Queue | 3 | 3 | 4 | 3 | 3 | 2 | 1 | 1 |
| DSCP | 61 | 53 | 45 | 37 | 29 | 21 | 13 | 5 |
| Queue | 3 | 3 | 4 | 3 | 3 | 2 | 1 | 1 |
| DSCP | 60 | 52 | 44 | 36(AF42) | 28(AF32) | 20(AF22) | 12(AF12) | 4 |
| Queue | 3 | 3 | 4 | 3 | 3 | 2 | 1 | 1 |
| DSCP | 59 | 51 | 43 | 35 | 27 | 19 | 11 | 3 |
| Queue | 3 | 3 | 4 | 3 | 3 | 2 | 1 | 1 |
| DSCP | 58 | 50 | 42 | 34(AF41) | 26(AF31) | 18(AF21) | 10(AF11) | 2 |
| Queue | 3 | 3 | 4 | 3 | 3 | 2 | 1 | 1 |
| DSCP | 57 | 49 | 41 | 33 | 25 | 17 | 9 | 1 |
| Queue | 3 | 3 | 4 | 3 | 3 | 2 | 1 | 1 |
| DSCP | 56(CS7) | 48(CS6) | 40(CS5) | 32(CS4) | 24(CS3) | 16(CS2) | 8(CS1) | 0(BE) |
| Queue | 3 | 3 | 4 | 3 | 3 | 2 | 1 | 1 |

The queue 4 is the highest queue and the default classes in the parentheses are defined by IETF.
To map DSCP to queues:

1. Click *Configuration > Quality of Service > DSCP to Queue*.



2. Select the Output Queue (traffic forwarding queue) to which the DSCP value is mapped.
3. Click **Apply**. 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:

1. Click *Configuration > Quality of Service > Bandwidth Control*

The *Bandwidth* page displays bandwidth information for each interface.

2. Select an interface, and click **Edit**.

Edit Bandwidth Control

Select Your Interface

Interface: Port GE13 LAG 1

Interface Settings

Ingress Rate Control: Enable

Ingress Rate Limit: 100 Kbits/sec (100-1000000)

Ingress Committed Burst Size: 128000 Bytes (3000-19173960)

Egress Shaping Control: Enable

Egress Committed Information Rate: 64 Kbits/sec (64-1000000)

Egress Committed Burst Size: 128000 Bytes (4096-16762902)

Apply Close

3. Select the Port or LAG interface.

4. Enter the 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.
- 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.

Note—The above *Ingress Rate Limit* fields do not appear when the interface type is LAG.

- 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 (CBS)—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.
5. Click **Apply**. The bandwidth settings are written to the Running Configuration file.

Egress Shaping

In addition to limiting transmission rate per port, which is done in the Bandwidth page, the device can limit the transmission rate of selected egressing frames on a per-queue per-port basis. Egress rate limiting is performed by shaping the output load.

The device limits all frames except for management frames. Any frames that are not limited are ignored in the rate calculations, meaning that their size is not included in the limit total.

Per-queue Egress rate shaping can be disabled.

To define egress shaping per queue:

1. Click *Configuration > Quality of Service > Egress Shaping*.

The Egress Shaping page displays the rate limit and burst size for each queue.

2. Select an interface type (Port or LAG), and click **Search**.
3. Select a Port/LAG, and click **Edit**.

Edit Egress Shaping

Select Your Interface

Interface: Port GE13 LAG 1

Egress Shaping Settings

Queue 1: Enable
Committed Information Rate: kbits/sec (64-1000000)
Committed Burst Size (CBS): Bytes (4096-16762902)

Queue 2: Enable
Committed Information Rate: kbits/sec (64-1000000)
Committed Burst Size (CBS): Bytes (4096-16762902)

Queue 3: Enable
Committed Information Rate: kbits/sec (64-1000000)
Committed Burst Size (CBS): Bytes (4096-16762902)

Queue 4: Enable
Committed Information Rate: kbits/sec (64-1000000)
Committed Burst Size (CBS): Bytes (4096-16762902)

Apply **Close**

This page enables shaping the egress for up to four queues on each interface.

4. Select the Interface.
5. For each queue that is required, enter the following fields:
 - Queue x—Select to enable egress shaping on queue number x.

- Committed Information Rate—Enter the maximum rate (CIR) in Kbits per second (Kbps). CIR is the average maximum amount of data that can be sent.
 - Committed Burst Size—Enter the maximum burst size (CBS) in bytes. CBS is the maximum burst of data allowed to be sent even if a burst exceeds CIR.
6. 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.

To configure Basic QoS:

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 use the DSCP value in the IP header.

In Basic QoS Mode, it is recommended that you disable the trusted mode at the ports where the CoS/802.1p and/or DSCP values of the incoming packets are not trustworthy. Otherwise, it might negatively affect the performance of your network. Incoming packets from ports that are disabled without trust mode are forwarded in best effort.

The Basic QoS page contains information for enabling Trust on the device. This configuration is only active when the QoS mode is Basic mode. Packets entering a QoS domain are classified at the edge of the QoS domain.

To define the Trust configuration:

Click *Configuration > Quality of Service > Basic QoS*.

Select the *Trusted Mode* while the device is in Basic mode. 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—All IP traffic is mapped to queues based on the values in their DSCP field. All non IP traffic is mapped to queues based on their CoS/802.1p value.

To disable QoS on a port:

1. Select a port or LAG , click **Edit**

2. Deselect QoS on the port or LAG.
3. Click **Apply**. The Running Configuration file is updated with the new settings.

Advanced QoS

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 ACLs of a class map with Permit (forward) action are considered belonging to the same flow. They 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 type 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 are 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.

Notes

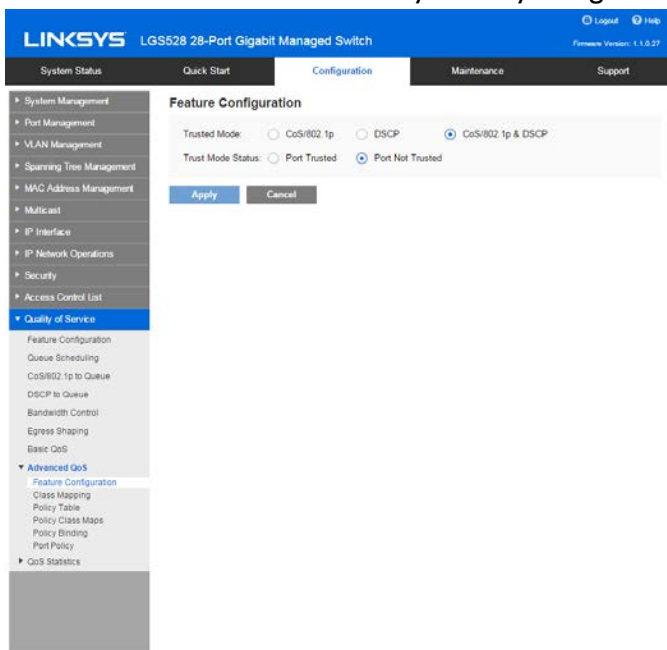
- *Single policer and aggregation policer are available when the device is in Layer 2 mode.*
- *An ACL can be configured to one or more class maps regardless of policies.*
- *A class map can belong to only one policy.*
- *When a class map using single policer is bound to multiple ports, each port has 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.

To configure Advanced QoS mode:

1. Select Advanced mode for the system by using the Feature Configuration page.



2. Select the Trust Mode from this page:
 - 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 on [p. 244](#).

If ACLs were defined, create class maps and associate the ACLs with them by using the *Class Mapping* page.
4. 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.
5. Bind the policy to an interface by using the *Policy Binding* page.

Feature Configuration (Advanced Mode)

The Basic QoS 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:

1. Click *Configuration > Quality of Service > Advanced QoS > Feature Configuration*.
2. Select the *Trust Mode* while the device is in Advanced 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), 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—All IP traffic is mapped to queues based on the values in their DSCP field. All non IP traffic is mapped to queues based on their CoS/802.1p value.
3. Select the default Advanced QoS trust mode (either Port Trusted or Port Not Trusted) for ports in the Trust Mode Status 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 Advanced QoS when the Trust Mode Status is set to Port Not Trusted, the default CoS values configured on the interface is ignored and all the traffic goes to queue 1. See the *Quality of Service > Advanced QoS > Global Settings* page for details.

If you have a policy on an interface then the Trust Mode Status is irrelevant, the action is according to the policy configuration.

Out-of-Profile DSCP Mapping

When a policer is assigned to a class maps (flows), you can specify the action to take when the amount of traffic in the flow(s) exceeds the

QoS-specified limits. The portion of the traffic that causes the flow to exceed its QoS limit is referred to as out-of-profile packets.

If the exceed action is Out of Profile DSCP, the device remaps the original DSCP value of the out-of-profile IP packets with a new value based on the Out of Profile DSCP Mapping Table. The device uses the new values to assign resources and the egress queues to these packets. The device also physically replaces the original DSCP value in the out of profile packets with the new DSCP value.

To use the out-of-profile DSCP exceed action, remap the DSCP value in the *Out Of Profile DSCP Mapping Table*. Otherwise the action is null, because the DSCP value in the table remaps the packets to itself by factory default.

This feature changes the DSCP tags for incoming traffic switched between trusted QoS domains. Changing the DSCP values used in one domain, sets the priority of that type of traffic to the DSCP value used in the other domain to identify the same type of traffic.

To remap DSCP values:

1. Click *Configuration > Quality of Service > Advanced QoS > Out of Profile DSCP Mapping*. This page enables setting the change-the-DSCP-value of traffic entering or leaving the device.

| | In DSCP | Out DSCP | In DSCP | Out DSCP | In DSCP | Out DSCP | In DSCP | Out DSCP |
|----|---------|----------|---------|----------|---------|----------|---------|----------|
| 0 | 0 | 16 | 16 | 32 | 32 | 48 | 48 | |
| 1 | 1 | 17 | 17 | 33 | 33 | 49 | 49 | |
| 2 | 2 | 18 | 18 | 34 | 34 | 50 | 50 | |
| 3 | 3 | 19 | 19 | 35 | 35 | 51 | 51 | |
| 4 | 4 | 20 | 20 | 36 | 36 | 52 | 52 | |
| 5 | 5 | 21 | 21 | 37 | 37 | 53 | 53 | |
| 6 | 6 | 22 | 22 | 38 | 38 | 54 | 54 | |
| 7 | 7 | 23 | 23 | 39 | 39 | 55 | 55 | |
| 8 | 8 | 24 | 24 | 40 | 40 | 56 | 56 | |
| 9 | 9 | 25 | 25 | 41 | 41 | 57 | 57 | |
| 10 | 10 | 26 | 26 | 42 | 42 | 58 | 58 | |
| 11 | 11 | 27 | 27 | 43 | 43 | 59 | 59 | |
| 12 | 12 | 28 | 28 | 44 | 44 | 60 | 60 | |
| 13 | 13 | 29 | 29 | 45 | 45 | 61 | 61 | |
| 14 | 14 | 30 | 30 | 46 | 46 | 62 | 62 | |
| 15 | 15 | 31 | 31 | 47 | 47 | 63 | 63 | |

2. DSCP In displays the DSCP value of the incoming packet that needs to be re-marked to an alternative value.
3. Select the DSCP Out value to where the incoming value is mapped.
4. Click **Apply**. The Running Configuration file is updated with the new DSCP Mapping table.

Class Mapping

A Class Map defines a traffic flow with ACLs (Access Control Lists). A MAC ACL, IPv4 ACL, and IPv6 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 has no 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 ACLs, beginning on [p. 243](#)).

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:

1. Click *Configuration > Quality of Service > Advanced QoS > Class Mapping*.

This page displays the already-defined class maps.

2. Click **Add**.

The screenshot shows the 'Add Class Map' configuration interface. It includes a text input for the class map name, radio buttons for selecting match criteria (IP, MAC, IP and MAC, IP or MAC), and dropdown menus for selecting specific ACLs for IP and MAC. There are also 'Apply' and 'Close' buttons at the bottom.

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 only one of the ACLs selected.

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:
 - 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.
4. Click **Apply**. The Running Configuration file is updated.

Aggregate Policers

You can measure the rate of traffic that matches a pre-defined 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. There are 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.

An aggregate policer applies the QoS to one or more class maps, therefore to 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.

To define an aggregate policer:

1. Click *Configuration > Quality of Service > Advanced QoS > Aggregate Policer*.

This page displays the existing aggregate policers.

2. Click **Add**.

Add Aggregate Policer

Enter New Policer

Aggregate Policer Name:

Policer Settings

Ingress Committed Information Rate: Kbits/sec (3-57982058)

Ingress Committed Burst Size: Bytes (3000-19173960)

Exceed Action: Forward Drop Out of Profile DSCP

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. Possible values:
 - Forward—Packets exceeding the defined CIR value are forwarded.
 - Drop—Packets exceeding the defined CIR value are dropped.
 - Out of Profile DSCP—The DSCP values of packets exceeding the defined CIR value are remapped to a value based on the Out Of Profile DSCP Mapping Table.
4. Click **Apply**. The Running Configuration file is updated.

Policy Table

The Policy Table page displays the list of advanced QoS policies defined in the system. The page also allows you to create and delete policies. Only those policies that are bound to an interface are active (see Policy Binding on [p. 277](#)).

Each policy consists of:

- 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.

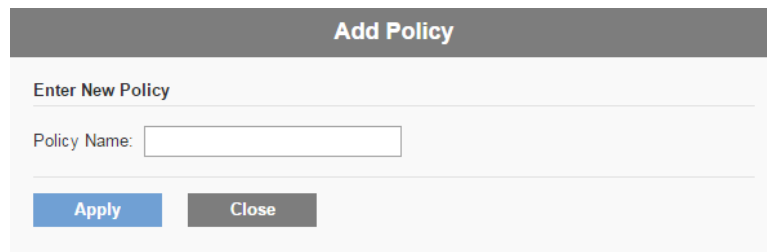
After a policy has been added, class maps can be added by using the Policy Table page.

To add a QoS policy:

1. Click *Configuration > Quality of Service > Advanced QoS > Policy Table*.

This page displays the list of defined policies.

2. Click **Add**.



3. Enter the name of the new policy in the *Policy Name* field.
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:

1. Click *Configuration > Quality of Service > Advanced QoS > Policy Class Maps*.
2. To add a new class map, click **Add**.
3. 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.
- 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 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..4) 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:
 - None—No policy is used.
 - Single Policer—The policer for the policy is a single policer.

If Police Type is Single Policer, enter the following QoS parameters:

 - o Ingress CIR—Enter the CIR in Kbps. See a description of this in the Bandwidth page.
 - o Ingress CBS—Enter the CBS in bytes. See a description of this in the Bandwidth page.
 - o Exceed Action—Select the action assigned to incoming packets exceeding the CIR. The options are:
 - o None—No action.
 - o Drop—Packets exceeding the defined CIR value are dropped.

- Out of Profile DSCP—IP packets exceeding the defined CIR are forwarding with a new DSCP derived from the Out Of Profile DSCP Mapping Table.
- Aggregate Policer—Select the policer for the policy is an aggregate policer.

4. 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:

1. Click *Configuration > Quality of Service > Advanced QoS > Policy Binding*.

The screenshot displays the 'Policy Binding' configuration page in the Linksys web interface. The page header shows 'LINKSYS LGS528 28-Port Gigabit Managed Switch' and 'Firmware Version: 1.1.0.27'. The navigation menu includes 'System Status', 'Quick Start', 'Configuration', 'Maintenance', and 'Support'. The left sidebar shows a tree view with 'Quality of Service' expanded to 'Policy Binding'. The main content area is titled 'Policy Binding' and contains a 'FILTER' section with 'Policy Name' and 'Interface Type' dropdowns, a 'Search' button, and a 'POLICY BINDING' table. The table has three sections for interfaces GE1-GE12, GE13-GE24, and GE25-GE28. Each section has a 'Bind' row and a 'Permit Any Unmatched Packets' row, each with checkboxes for every interface. 'Apply' and 'Cancel' buttons are at the bottom.

2. Select a Policy Name and Interface Type if required.
3. Click **Search**.
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.*

5. Click **Apply**. The QoS policy binding is defined, and the Running Configuration file is updated.

Port Policy

To view a policy on a port:

1. Click *Configuration > Quality of Service > Advanced QoS > Port Policy*.
2. Select an Interface Type and a Policy Name.
3. Click Search. The policy is selected.
4. The following fields are displayed for the port and policy selected:
 - Interface—Name of interface.
 - Policy Name—Name of policy.
 - Permit Any Unmatched Packets—Whether the feature of forwarding 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.*

5. Click **Apply**. The QoS policy binding is defined, and the Running Configuration file is updated.

QoS Statistics

Queue Statistics

The *Queue Statistics* page displays queue statistics, including statistics of forwarded and dropped packets, based on interface, queue, and drop precedence.

To view Queue Statistics:

1. Click *Configuration > Quality of Service > QoS Statistics > Queues Statistics*.

This page displays the following fields:

- Refresh Rate—Select the time period that passes before the interface Ethernet statistics are refreshed. The available options:
 - 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.
 - Counter Set—The options are:
 - Set 1—Displays the statistics for Set 1 that contains all interfaces and queues with a high DP (Drop Precedence).
 - Set 2—Displays the statistics for Set 2 that contains all interfaces and queues with a low DP.
 - Interface—Queue statistics are displayed for this interface.
 - Queue—Packets were forwarded or tail dropped from this queue.
 - Drop Precedence—Lowest drop precedence has the lowest probability of being dropped.
 - Total Packets—Number of packets forwarded or tail dropped.
 - Tail Drop Packets—Percentage of packets that were tail dropped.
2. Click **Add**.
 3. Enter the parameters.
 - Counter Set—Select the counter set:
 - Set 1—Displays the statistics for Set 1 that contains all interfaces and queues with a high DP (Drop Precedence).
 - Set 2—Displays the statistics for Set 2 that contains all interfaces and queues with a low DP.

- Interface—Select the ports for which statistics are displayed. The options are:
 - Port—Selects the port on the selected unit number for which statistics are displayed.
 - All Ports—Specifies that statistics are displayed for all ports.
 - Queue—Select the queue for which statistics are displayed.
 - Drop Precedence—Enter drop precedence that indicates the probability of being dropped.
4. Click **Apply**. The Queue Statistics counter is added, and the Running Configuration file is updated.

Chapter 15 – Maintenance

This section describes how to view system information and configure various options on the device.

It covers the following topics:

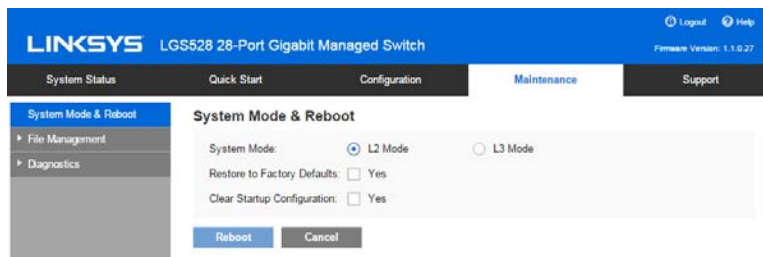
- [Device Models](#)
- [System Mode & Reboot](#)
- [File Management](#)
- [Diagnostics](#)

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

Restore Factory Defaults

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.

You can back up the device configuration by using *Maintenance > File Management > Configuration File Copy*. You can also upload the configuration from a remote device.

System Modes

The device can be in either Layer 2 or Layer 3 system mode.

Note—If you change the system mode after clicking *Apply*, the system will require a reboot, and the startup configuration file will be removed after the boot.

To configure the system mode or reboot to factory defaults:

1. Click *Maintenance > System Mode & Reboot*.
2. Enter the following fields:
 - System Mode—Select either Layer 2 or Layer 3 system mode.
 - Restore to Factory Defaults—Select to reboot the device by using the factory default configuration. This process erases the Startup Configuration file and the backup configuration file.
 - Clear Startup Configuration—Select to clear the startup configuration on reboot.
3. Click **Reboot**. The parameters are copied to the Running Configuration file and the stack is rebooted.

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](#)
- [DHCP Auto Configuration](#)

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.
- TFTP 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.
- Enable automatically uploading a configuration file from a DHCP server to the device, as described in the Auto Configuration via DHCP 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

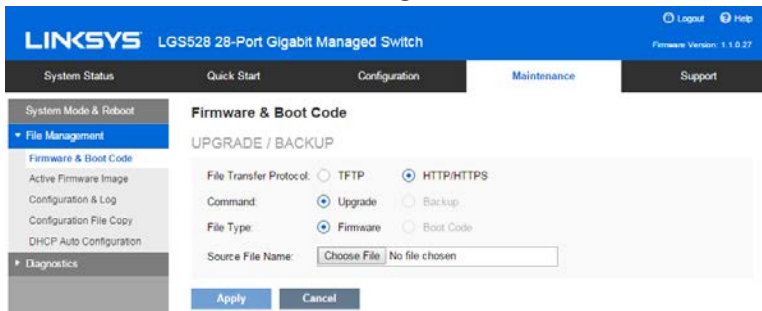
There are two firmware images stored on the device. One of the images is identified as the active image and other image is identified as the inactive image.

When you upgrade the firmware, the new image always replaces the image identified as the inactive image.

Even after uploading new firmware on the device, the device continues to boot by using the active image (the old version) until you change the status of the new image to be the active image by using the procedure in the Active Firmware Image section. Then boot the device.

To upgrade or backup a software image:

1. Click *Maintenance > File Management > Firmware & Boot Code*.



The screenshot shows the Linksys web interface for an LGS528 28-Port Gigabit Managed Switch. The page title is "LINKSYS LGS528 28-Port Gigabit Managed Switch" and the firmware version is "1.1.0.27". The navigation menu includes "System Status", "Quick Start", "Configuration", "Maintenance", and "Support". The "Maintenance" menu is expanded, showing "System Mode & Reboot", "File Management", "Firmware & Boot Code", "Active Firmware Image", "Configuration & Log", "Configuration File Copy", "DHCP Auto Configuration", and "Diagnostics". The "File Management" menu is also expanded, showing "Firmware & Boot Code". The "Firmware & Boot Code" page has a sub-header "UPGRADE / BACKUP". The "File Transfer Protocol" is set to "HTTP/HTTPS". The "Command" is set to "Upgrade". The "File Type" is set to "Firmware". The "Source File Name" field is empty, with a "Choose File" button and "No file chosen" text. There are "Apply" and "Cancel" buttons at the bottom.

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.
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.
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.
5. Click **Apply**.

Note—When the process is completed, the following information is displayed:

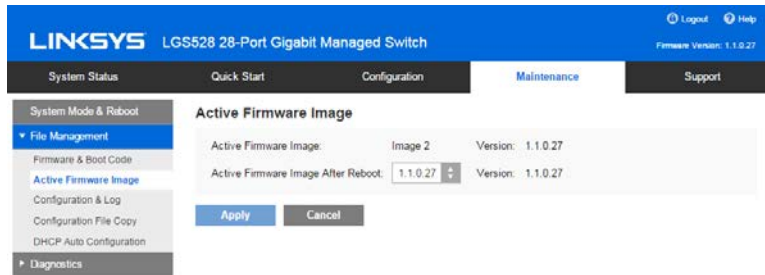
- *Bytes Transferred*—How many bites were transferred in the process.
- *Status*—Did the process succeed or fail.
- *Error Message*—Reason for failure of the process.

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:

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.
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.
 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.

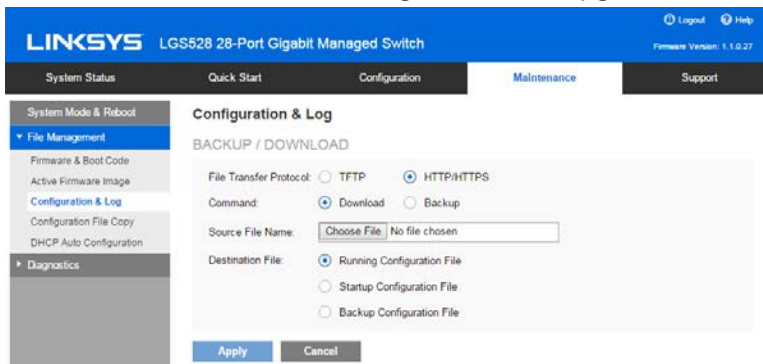
When restoring a configuration file to the Running Configuration, the imported file adds any configuration commands that did not exist in the old file and overwrites any parameter values in the existing configuration commands.

When restoring a configuration file to the 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:

1. Click *Maintenance > File Management > Configuration & Log*.



2. Select the File Transfer Protocol.
3. If you selected via TFTP, enter the parameters. Otherwise, skip to STEP 4.

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.*

4. Click **Apply**. The file is upgraded or backed up.
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.
6. Only valid file types are displayed. (The file types are described in the Files and File Types section).
7. 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:

1. Click *Maintenance > File Management > Configuration File Copy*.
2. Select the Source File to be copied. Only valid file types are displayed (described in the Files and File Types section).
3. Select the Destination File to be overwritten by the source file.
4. Click **Apply**. The file is copied.

DHCP Auto Configuration

The Auto Configuration Update feature provides a convenient method to automatically configure devices in a network. This process enables the administrator to remotely ensure that the configuration of these devices in the network is up-to-date.

DHCP Auto Configuration downloads a configuration file from a remote TFTP server. At the end of the Auto Configuration process, the device reboots itself using the configuration file.

To use this feature, configure a DHCP server in the network with the locations and names of the configuration files of your devices. The devices in the network are configured as DHCP clients by default. When the devices are assigned their IP addresses by the DHCP server, they also receive information about the configuration file. If the configuration file is different from the one currently used on the device, the device reboots itself after downloading the file.

Auto Configuration Process

DHCP Auto Configuration uses the configuration server name/address and configuration file name/path, if any, in the DHCP messages received. This information is specified as DHCP options in the Offer message coming from the DHCPv4 servers and in the Information Reply messages coming from DHCPv6 servers.

If this information is not found in the DHCP server messages, backup information that has been configured in the DHCP Auto Configuration page is used.

When the Auto Configuration process is triggered (see Auto Configuration Update Trigger), the sequence of events described below occurs.

- The device uses the TFTP server name/address and configuration file name/path (DHCPv4 options: 66,150, and 67, DHCPv6 options: 59 and 60), if any, from the DHCP message received.
- If the information is not sent by the DHCP server, the Backup Server IP Address/Name and the Backup Configuration File Name (from the DHCP Auto Configuration page) is used.
- The new configuration file is used if its name is different than the name of the configuration file previously used on the device or if the device has never been configured.
- The device is rebooted with the new configuration file, at the end of the Auto Configuration Process.
- SYSLOG messages are generated by the copy process.

Auto Configuration Update Trigger

Auto Configuration via DHCPv4 is triggered when the following conditions are fulfilled:

- The IP address of an IP interface is configured as dynamic, and is assigned by a DHCP server at reboot, explicit administration renewal action, or automatic renewal of an expiring lease. Explicit renewal can be activated in the IPv4 Interface page.
- If Auto Configuration is enabled, the Auto Configuration process is triggered when the configuration file name is received from a DHCP server or a backup configuration file name has been configured.

Auto Configuration via DHCPv6 is triggered when the following conditions are fulfilled:

- When a DHCPv6 server sends information to the device, as in the following cases:
 - When an IPv6-enabled interface is defined as a DHCPv6 stateless configuration client.
 - When DHCPv6 messages are received from the server,
 - When DHCPv6 information is refreshed by the device.
 - After rebooting the device when stateless DHCPv6 client is enabled.

- When the DHCPv6 server packets contain the configuration filename option.

Ensuring Correct Performance

To ensure that the DHCP Auto Configuration feature works correctly, note the following:

- A configuration file that is placed on the TFTP server must match the form and format requirements of the supported configuration file. The form and format of the file are checked, but the validity of the configuration parameters is not checked prior to loading it to the Startup Configuration.
- In IPv4, to ensure that a device downloads the configuration file as intended during the Auto Configuration process, it is recommended that the device is always assigned the same IP address. This ensures that the device is always assigned with the same IP address, and obtains the same information used in Auto Configuration.

Default Settings and Configuration

The following defaults exist on the system:

- Auto Configuration is enabled.
- The device is enabled as a DHCP client.

Before You Start the Auto Configuration Process

To use this feature, the device must either be configured as a DHCPv4 or DHCPv6 client. The type of DHCP client defined on the device is in correlation with the type of interfaces defined on the device. To do this, set the address type of the interface to Dynamic (in the IPv4 Interface and IPv6 Interface pages).

Auto Configuration Preparations on the Server

To prepare the DHCP and TFTP servers:

TFTP Server

- Place a configuration file in the working directory. This file can be created by copying a configuration file from a device. When the device is booted, this becomes the Running Configuration file.

DHCP Server

- DHCPv4:
 - 66 (single server address) or 150 (list of server addresses)
 - 67 (name of configuration file)

- DHCPv6
 - Option 59 (server address)
 - Options 60 (name of configuration file)
1. Configure Auto Configuration parameters in the *Maintenance > File Management > DHCP Auto Configuration* page.
 2. Set the IP Address Type to Dynamic in the *Configuration > IP Interface > IPv4 > IPv4 Interface* page.
 3. Enable IPv6 Address Auto Configuration in the *Configuration > IP Interface > IPv6 > IPv6 Interface* page.

To configure DHCP Auto Configuration:

1. Click *Maintenance > File Management > DHCP Auto Configuration*.
2. Enter the values.
 - Auto Configuration Via DHCP—Select this field to enable DHCP Auto Configuration. This feature is enabled by default, but can be disabled here.
 - Backup Server—Select whether the backup server will be configured *By Address* or *By 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:
 - Global—The IPv6 address is a global Unicast IPV6 type that is visible and reachable from other networks. Select the link local interface (if IPv6 is used) from the list.
 - 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 is used) from the list.
3. Enter the following optional information that is used if the DHCP server did not provide the required information.
 - Backup Server IP Address—Enter the backup server IP address.
 - Backup Server Name—Enter the backup server name.
 - Backup Configuration File Name—Enter the backup configuration file name.

The following fields are displayed:

- Last Auto Configuration Server—Address of the last backup server.
 - Last Auto Configuration File Name—Name of the last configuration file name.
4. Click **Apply**. The parameters are copied to the Running Configuration file.

Diagnostics

This section covers the following topics:

- [Copper Test](#)
- [Optical Module Status](#)
- [Ping](#)
- [Traceroute](#)
- [Port Mirroring](#)

Copper Test

The Copper Test page displays the results of integrated cable tests performed on copper cables by the Virtual Cable Tester (VCT).

VCT performs two types of tests:

- Time Domain Reflectometry (TDR) technology tests the quality and characteristics of a copper cable attached to a port. Cables of up to 140 meters long can be tested. These results are displayed in the Test Results block of the Copper Test page.
- DSP-based tests are performed on active GE links to measure cable length.

These results are displayed in the Advanced Information block of the Copper Test page.

Preconditions to Running the Copper Port Test

Before running the test:

- (Mandatory) Disable Short Reach mode (see the Configuration > Port Management > Green Ethernet > Properties page)
- (Optional) Disable EEE (see the Configuration > Port Management > Green Ethernet > Properties page)

Use a CAT5 data cable when testing cables using (VCT).

Accuracy of the test results can have an error range of +/- 10 for Advanced Testing and +/- 2 for basic testing.

Caution—When a port is tested, it is set to the Down state and communications are interrupted. After the test, the port returns to the Up state. It is not recommended that you run the copper port test on a port you are using to run the web-based switch configuration utility, because communications with that device are disrupted.

To test copper cables attached to ports:

1. Click Maintenance > Diagnostics > Copper Test.
2. Select the port on which to run the test.
3. Click Test.
4. When the message appears, click OK to confirm that the link can go down or Cancel to abort the test.

The following fields are displayed in the Test Results block:

- Test Results—Cable test results. Possible values are:
 - OK—Cable passed the test.
 - No Cable—Cable is not connected to the port.
 - Open Cable—Cable is connected on only one side.
 - Short Cable—Short circuit has occurred in the cable.
 - Unknown Test Result—Error has occurred.
- Distance to Fault—Distance from the port to the location on the cable where the fault was discovered.
- Port Operational Status—Displays whether port is up or down.

Note—TDR tests cannot be performed when the port speed is 10Mbit/Sec.

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*.

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:

1. Click *Maintenance > Diagnostics > Ping*.

2. Configure ping by entering the fields:

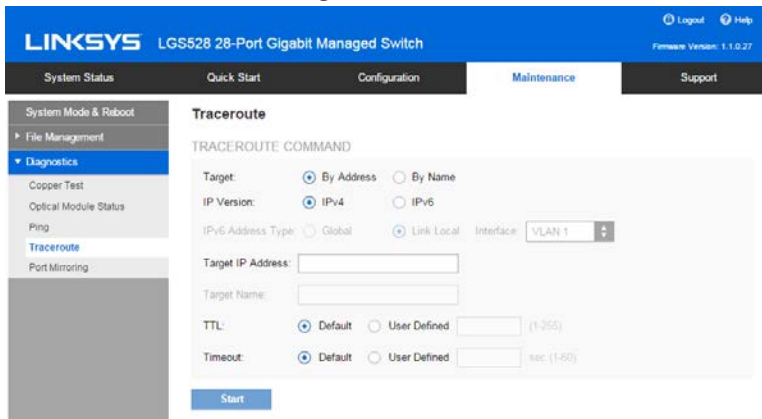
- Target—Select whether to specify the source interface by its IP address or name. This field influences the interfaces that are displayed in the Source IP field, as described below.

- 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.
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.
 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

Traceroute

Traceroute discovers the IP routes along which packets were forwarded by sending an IP packet to the target host and back to the device. The Traceroute page shows each hop between the device and a target host, and the round-trip time to each such hop.

1. Click *Maintenance > Diagnostics > Traceroute*.



The screenshot shows the Linksys web interface for an LGS528 28-Port Gigabit Managed Switch. The page title is "LINKSYS LGS528 28-Port Gigabit Managed Switch" and the firmware version is "1.1.0.27". The navigation menu includes System Status, Quick Start, Configuration, Maintenance, and Support. The left sidebar shows System Mode & Reboot, File Management, Diagnostics (selected), Copper Test, Optical Module Status, Ping, Traceroute (selected), and Port Mirroring. The main content area is titled "Traceroute" and contains the "TRACEROUTE COMMAND" configuration form. The form includes radio buttons for "By Address" (selected) and "By Name", radio buttons for "IPv4" (selected) and "IPv6", radio buttons for "Global" and "Link Local" (selected), and a dropdown menu for "Interface" set to "VLAN 1". There are input fields for "Target IP Address", "Target Name", "TTL" (with "Default" selected), and "Timeout" (with "Default" selected). A "Start" button is at the bottom.

2. Configure Traceroute by entering information into the following fields:
 - Target—Select whether target hosts are identified by their IP address or name.
 - IP Version—If the target host 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 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. If this mode
 - Global—The IPv6 address is a global Unicast IPV6 type that is visible and reachable from other networks. Select the link local interface (if IPv6 is used) from the list.
 - Target IP Address—Select the target source interface whose IPv4 address will be used as the source IPv4 address for communication messages. Only the existing IP addresses of the type specified in the IP Version field will be displayed.
 - Target Name—Enter the target host name.

- TTL—Enter the maximum number of hops that Traceroute permits. This is used to prevent a case where the sent frame gets into an endless loop. The Traceroute command terminates when the destination is reached or when this value is reached. To use the default value (30), select Use Default.
 - Timeout—Enter the length of time that the system waits for a frame to return before declaring it lost, or select Use Default.
3. Click **Start**. The operation is performed.

A page appears showing the Round Trip Time (RTT) and status for each trip in free text containing the following information:

- Index—Displays the number of the hop.
- Host—Displays a stop along the route to the destination.
- Round Trip Time (1-3)—Displays the round trip time in (ms) for the first through third frame and the status of the first through third operation.

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:

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.
2. Click **Add** to add a port to be mirrored.
 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.
 4. Click **Apply**. Port mirroring is added to the Running Configuration.

Chapter 16 - Support

Click **Get Support** to go to the Linksys Small Business support website. Resources available there include setup help, frequently asked questions, software downloads, live chat with technical support, and community forums.

Appendix

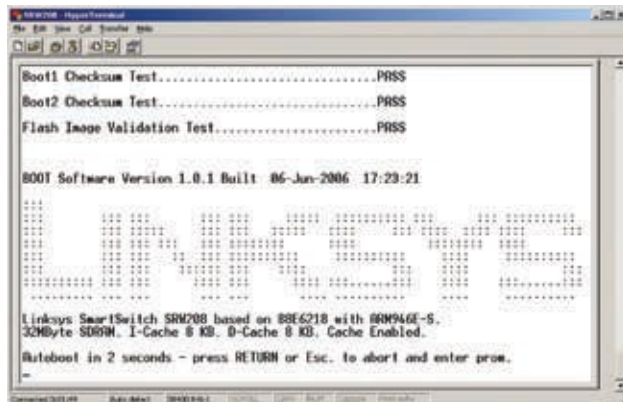
Startup Menu Procedures

The Startup menu can be entered when booting the device. There is a two second window of time to enter the Startup Menu immediately after the POST test. The menu can be accessed directly from a terminal connected to the console port. The Startup menu procedures can be done using the ASCII terminal or Windows HyperTerminal.

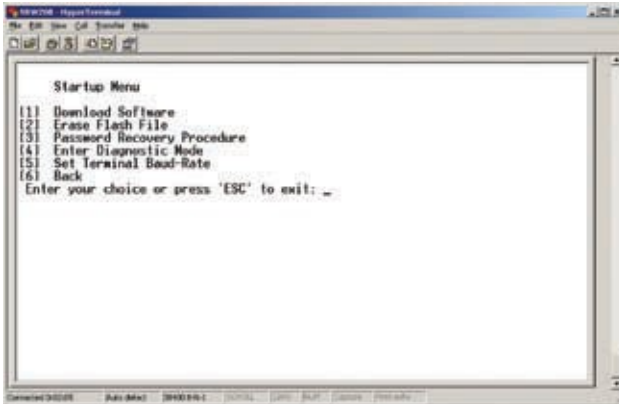
The software download procedure is performed when a new version must be downloaded to replace corrupted files, update or upgrade the system software. To download software from the Startup menu:

Enter the Startup menu:

1. Power off your computer and switch.
2. Connect the provided null modem cable from the COM port on your computer to the Console port on the switch.
3. Power on your computer and launch HyperTerminal, follow the instructions in Chapter 4: *Configuration Using the Console Interface to configure HyperTerminal* to connect to the switch.
4. Power on the switch and watch for the auto-boot message:
Autoboot in 2 seconds - press Return or Esc to abort and enter prom.

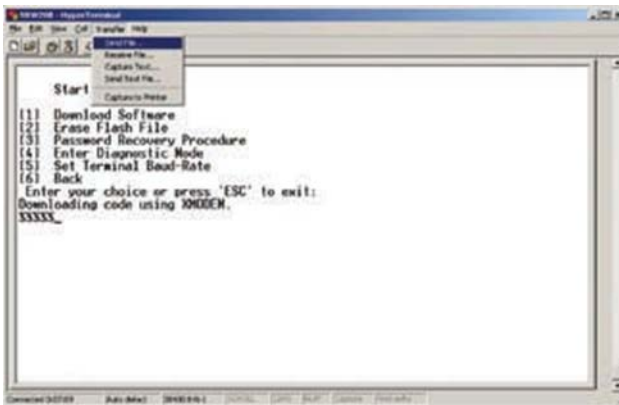


- When the auto-boot message appears, press the Enter key to access the Startup menu.



NOTE— If a selection is not made within 35 seconds (default), the device times out and you will need to disconnect the power to restart the process.

- Select **[1] Download Software** and a message will appear, downloading code using XMODEM with characters running across the screen. If you do not perform the steps on the next page to locate the file for download within a certain time, the device will reset.
- Select **Send File** from the *Transfer* drop-down menu.



- In the Filename field, enter the file path for the file to be downloaded or click Browse to locate the file.

Only valid files, with a *.ros or *.rfb suffix, that have been provided by Linksys, can be downloaded.

Downloading invalid files will result in unpredictable behavior.

Ensure that the Xmodem protocol is selected in the *Protocol*: field.

9. Press **Send** and the software is downloaded.



After the software has been downloaded, the device will reboot automatically.

Notes:

For regulatory, warranty, and safety information, see the CD that came with your switch or go to Linksys.com/support/.

Specifications are subject to change without notice.

Visit linksys.com/support/ for award-winning technical support.

BELKIN, LINKSYS and many product names and logos are trademarks of the Belkin group of companies. Third-party trademarks mentioned are the property of their respective owners.

© 2015 Belkin International, Inc. and/or its affiliates. All rights reserved.

LNKPG-00144 Rev. B00