



EdgeSwitch[™]

User Interface for PoE Switches
Models: ES-24-250W, ES-24-500W,
ES-48-500W, ES-48-750W

Administration Guide

Table of Contents

About This Document	7
Purpose and Audience	7
Document Organization	7
Products and Models	7
Related Documents	8
Typographical Conventions	8
Chapter 1: Getting Started	9
Connecting the Switch to the Network	9
Understanding the User Interfaces	9
Using the EdgeSwitch UI	9
Accessing the UI	9
EdgeSwitch UI Page Layout	10
Device View	10
Navigation Menu	11
Configuration and Status Fields	12
Table Filtering	14
Help Page Access	14
User-Defined Fields	14
Using the Command-Line Interface	15
Chapter 2: Configuring Power over Ethernet	16
Chapter 3: Configuring System Information	18
Viewing ARP Cache	19
Viewing Inventory Information	20
Viewing the Dual Image Status	21
Viewing System Resources	22
System Resource Status	22
System Resource Configuration	23
Defining General Device Information	24
System Description	25
IP Address Conflict Detection	26
Network Port IPv6 Neighbors	29
DHCP Client Options	30
Secure HTTP Configuration	31
SSH Configuration	32
Authentication Server Users	36
Logged in Sessions	38

Accounting Selection	40
Authentication Selection	43
Last Password Result	47
Denial of Service Configuration	48
CLI Banner Configuration	50
Basic Switch Configuration	51
Switch Configuration	51
Managing Logs	52
Log Configuration	52
Buffered Log	54
Event Log	55
Logging Hosts	56
Syslog Source Interface Configuration	57
Persistent Log	58
Configuring Email Alerts	59
Email Alert Global Configuration	59
Email Alert Server Configuration	60
Email Alert Statistics	61
Email Alert Subject Configuration	62
Email Alert To Address Configuration	63
Viewing Device Port Information	64
Port Summary	64
Port Description	66
Cable Test	67
Mirroring	68
Configuring a Port Mirroring Session	69
Configuring Port Mirroring Source Ports	69
Configuring the Port Mirroring Destination	70
Defining SNMP Parameters	71
SNMP v1 and v2	71
SNMP v3	71
SNMP Community Configuration	72
SNMP v1/v2 Trap Receivers Configuration	73
SNMP v3 Trap Receivers Configuration	74
SNMP Access Control Group	75
SNMP User Security Model	76
SNMP Trap Source Interface Configuration	77
Viewing System Statistics	79
Switch Detailed Statistics	79
Port Summary	81
Port Detailed Statistics	82
Network Port DHCPv6 Client Statistics	85

Time-Based Group Statistics	86
Time-Based Flow Statistics	87
Time-Based Statistics	89
Using System Utilities.....	90
System Reset.....	90
Ping	90
TraceRoute.....	93
IP Address Conflict Detection.....	94
Uploading Files	96
Downloading Files	97
AutoInstall	98
Managing SNMP Traps.....	100
System Trap Log.....	100
System Trap Flags	101
Managing the DHCP Server.....	102
DHCP Server Global Configuration.....	102
DHCP Server Pool Configuration.....	103
DHCP Server Pool Options.....	104
DHCP Server Bindings Information.....	106
DHCP Server Statistics.....	107
DHCP Server Conflicts Information.....	108
Configuring Time Ranges.....	109
Time Range Configuration.....	109
Time Range Entry Configuration.....	110
Configuring DNS	112
DNS Global Configuration	112
DNS IP Mapping Configuration	113
DNS Source Interface Configuration	114
Configuring SNTP Settings.....	115
SNTP Global Configuration	116
SNTP Global Status	117
SNTP Server Configuration	118
SNTP Server Status	119
Configuring the Time Zone	121
Time Zone Configuration.....	122
Summer Time Configuration.....	123
Chapter 4: Configuring Switching Information	125
Managing VLANs.....	126
VLAN Status.....	126
VLAN Port Configuration	127
VLAN Port Summary.....	128

VLAN Internal Usage	129
Reset VLAN Configuration	130
Managing Voice VLANs	131
Voice VLAN Configuration	131
Voice VLAN Interface Summary	131
Creating MAC Filters	133
MAC Filter Configuration	133
GARP Configuration	134
GARP Switch Configuration	134
Configuring DHCP Snooping	136
Global DHCP Snooping Configuration	136
DHCP Snooping Static Bindings	139
DHCP Snooping Dynamic Bindings	140
DHCP Snooping Persistent Configuration	141
DHCP Snooping Statistics	142
Configuring IGMP Snooping	143
Global Configuration and Status	143
Interface Configuration	144
IGMP Snooping Source Specific Multicast	145
IGMP Snooping VLAN Status	146
IGMP Snooping Multicast Router Configuration	147
IGMP Snooping Multicast Router VLAN Status	148
IGMP Snooping Multicast Router VLAN Configuration	149
Configuring IGMP Snooping Querier	150
IGMP Snooping Querier Configuration	150
VLAN Configuration	151
IGMP Snooping Querier VLAN Status	152
Creating Port Channels	153
Port Channel Summary	153
Port Channel Statistics	155
Viewing Multicast Forwarding Database Information	156
Multicast Forwarding Database Summary	156
Multicast Forwarding Database GMRP Table	157
Configuring Protected Ports	159
Configuring Spanning Tree Protocol	160
Spanning Tree Switch Configuration	160
Spanning Tree CST Configuration	161
Spanning Tree CST Port Configuration	163
Spanning Tree MST Configuration	165
Spanning Tree MST Port Configuration	167
Spanning Tree Statistics	169

Configuring Port Security	171
Port Security Global Administration	171
Port Security Interface Status	172
Port Security Statically Configured MAC Addresses	173
Port Security Dynamically Learned MAC Addresses	174
Managing LLDP	175
LLDP Global Configuration	175
LLDP Interface Configuration	176
LLDP Local Device Summary	177
Remote Device Summary	178
LLDP Statistics	179
LLDP-MED	181
LLDP-MED Global Configuration	181
LLDP-MED Local Device Information	183
LLDP-MED Remote Device Information	184
Chapter 5: Configuring Routing	186
Configuring ARP	187
ARP Table	188
ARP Table Configuration	189
Configuring Global IP Settings	190
Routing IP Configuration	190
Routing IP Interface Summary	192
Routing IP Interface Configuration	194
Routing IP Statistics	196
Router	199
Route Table	199
Configured Routes	200
Adding a Static Route	200
Configuring Policy-Based Routing	202
Chapter 6: Managing Device Security	203
Port Access Control	203
Global Port Access Control Configuration	204
Port Access Control Port Summary	205
Port Access Control Port Configuration	207
Port Access Control Port Details	210
Port Access Control Statistics	212
Port Access Control Client Summary	214
Port Access Control Privileges Summary	215
Port Access Control History Log Summary	216

RADIUS Settings	217
RADIUS Configuration	217
RADIUS Named Server Status	218
RADIUS Server Statistics	219
RADIUS Accounting Server Status	220
RADIUS Accounting Server Statistics	221
RADIUS Clear Statistics	222
RADIUS Source Interface Configuration	223
TACACS+ Settings	224
TACACS+ Configuration	224
TACACS+ Server Summary	225
TACACS+ Server Configuration	226
TACACS+ Source Interface Configuration	227
Chapter 7: Configuring Quality of Service	228
Configuring Access Control Lists	229
IP Access Control Lists	229
Access Control List Summary	230
Access Control List Configuration	231
Access Control List Interface Summary	235
Access Control List VLAN Summary	236
Configuring Auto VoIP	237
Auto VoIP Global Configuration	237
OUI Table Summary	238
OUI Based Auto VoIP	239
Protocol Based Auto VoIP	240
Configuring Class of Service	242
CoS IP DSCP Mapping Configuration	242
CoS Interface Queue Configuration	245
CoS Interface Queue Drop Precedence Configuration	246
Configuring Diffserv	247
Diffserv Global Configuration and Status	247
Diffserv Class Summary	248
Diffserv Class Configuration	249
Diffserv Policy Summary	252
Diffserv Policy Configuration	253
Diffserv Service Summary	255
Diffserv Service Performance Statistics	256
Diffserv Policy Performance Statistics	257

Appendix A: Configuration Examples	258
Configuring VLANs	258
Using the EdgeSwitch UI to Configure VLANs	258
Using the CLI to Configure VLANs.....	260
Configuring Multiple Spanning Tree Protocol	261
Using the Web UI to Configure MSTP.....	261
Using the CLI to Configure MSTP.....	263
Configuring VLAN Routing.....	264
Using the CLI to Configure VLAN Routing	264
Configuring Policy-Based Routing	266
Configuring Policy-Based Routing Using the CLI.....	266
Configuring 802.1X Network Access Control	269
Using the CLI to Configure 802.1X Port-Based Access Control.....	269
Configuring Differentiated Services for VoIP.....	270
Using the CLI to Configure DiffServ VoIP Support	270
Appendix B: Contact Information	272
Ubiquiti Networks Support	272
Online Resources.....	272

About This Document

This section contains the following information about this document:

- **“Purpose and Audience” on page 8**
- **“Document Organization” on page 8**
- **“Products and Models” on page 8**
- **“Related Documents” on page 9**
- **“Typographical Conventions” on page 9**
- **“Typographical Conventions” on page 9**

Purpose and Audience

This guide describes how to configure the EdgeSwitch software features using the browser-based EdgeSwitch user interface (UI). The information in this guide is intended for system administrators who are responsible for configuring and operating a network using EdgeSwitch devices.

To obtain the greatest benefit from this guide, you should have an understanding of the base software and should have read the specification for your networking device platform. You should also have basic knowledge of Ethernet and networking concepts.

Document Organization

This guide contains the following sections:

- **“Chapter 1: Getting Started” on page 10** contains information about performing the initial system configuration and accessing the user interface.
- **“Chapter 3: Configuring System Information” on page 19** describes how to configure administrative features such as SNMP, system users, and port information.
- **“Chapter 4: Configuring Switching Information” on page 126** describes how to manage and monitor the Layer-2 switching features.
- **“Chapter 5: Configuring Routing” on page 187** describes how to configure the Layer-3 routing features.
- **“Chapter 6: Managing Device Security” on page 204** contains information about configuring switch security information such as port access control, TACACS+, and RADIUS server settings.
- **“Chapter 7: Configuring Quality of Service” on page 229** describes how to manage the EdgeSwitch software ACLs, and how to configure the Differentiated Services and Class of Service features.
- **“Appendix A: Configuration Examples” on page 259** describes how to configure selected features on the switch using either the EdgeSwitch UI, command-line interface, and/or Simple Network Management Protocol (SNMP).

Products and Models

This document covers the following Ubiquiti products and models:

Affected Products

Name	Description	Part Number
EdgeSwitch 48-port 750W	Managed PoE+ Gigabit Switch with SFP+	ES-48-750W
EdgeSwitch 48-port 500W	Managed PoE+ Gigabit Switch with SFP+	ES-48-500W
EdgeSwitch 24-port 500W	Managed PoE+ Gigabit Switch with SFP	ES-24-500W
EdgeSwitch 24-port 250W	Managed PoE+ Gigabit Switch with SFP	ES-24-250W

Related Documents

- *EdgeSwitch CLI Command Reference*
- *EdgeSwitch Quick Start Guide*

For additional information, refer to the EdgeSwitch community website: community.ubnt.com/edgemax

Typographical Conventions

The following table lists typographical conventions used throughout this document.

Typographical Conventions

Convention	Indicates	Example
Bold	User selection User-entered text	Select VLAN 2 from the <i>VLAN ID</i> list; Click Submit enter 3 to assign VLAN 3 as the default VLAN
<i>Italic</i>	Name of a field Name of UI page, dialog box, window, etc.	delete the existing name in the <i>Username</i> field Use the <i>IP Address Conflict Detection</i> page
>	Order of navigation selections to access a page	To access the <i>Session</i> page, click System > Users > Session
Courier font	CLI commands and their output	show network

Chapter 1: Getting Started

This chapter describes how to start the switch and access the user interface. It contains the following sections:

- **[“Connecting the Switch to the Network” on page 10](#)**
- **[“Understanding the User Interfaces” on page 10](#)**

Connecting the Switch to the Network

For detailed instructions on how to connect the EdgeSwitch to your network, refer to the *Quick Start Guide* that came with the switch.

Understanding the User Interfaces

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

- EdgeSwitch User Interface (UI)
- Command-Line Interface (CLI)

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

This guide describes how to use the EdgeSwitch UI to manage and monitor the system. For information about how to manage and monitor the system by using the CLI, see the *EdgeSwitch CLI Command Reference*.

Using the EdgeSwitch UI

This section describes how to use the EdgeSwitch UI.

Accessing the UI

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

- HTML version 4.0, or later
- HTTP version 1.1, or later
- JavaScript® version 1.5, or later

Use the following procedures to log into the EdgeSwitch UI:

1. Open a web browser and enter the IP address of the switch in the web browser address field. The login screen appears, as shown in the following illustration.



EdgeSwitch UI Login Screen


2. Type the *User Name* and *Password* into the fields on the login screen, and then click **Login**.

The user name and password are the same as those you use to log on to the command-line interface. By default, the user name is *ubnt*, and the password is *ubnt*. Passwords are case-sensitive.

3. If this is your first login to the UI, read the license agreement. Then, click the **I agree to the terms of this License Agreement** check box and click **Log In**.
4. After the system authenticates you, the *System Description* page is displayed.

EdgeSwitch UI Page Layout

The following illustration shows the layout of a page in the EdgeSwitch UI. Each UI page contains three main areas: the device view, the navigation menu, and the configuration and status fields. Each page also provides buttons that let you perform operations on the displayed information, access a context-specific help page, or log out of the system.



EdgeSwitch UI Page Layout

Device View

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



Example of Device View (24-Port Models)

In the Device View, colors indicate status information:

- Gray indicates that the port link is down.
- Amber indicates that the port link is up at 100 Mbps.
- Green indicates that the port link is up at 1 Gbps.
- A white dot indicates PoE output.

To obtain additional information on the ports:

- Click any port to display the *Port Description* page (**System > Port > Description** on the navigation menu).
- Right-click any port to display a menu with links to the following port configuration-related pages:
 - *PoE* (**PoE > PoE Configuration**)
 - *Port Summary* (**System > Port > Summary**)
 - *Port Description* (**System > Port > Description**)
 - *Port Cable Test* (**System > Port > Cable Test**)
 - *Multiple Port Mirroring* (**System Port Mirroring**)
 - *Port Summary Statistics* (**System > Statistics > System > Port Summary**)
 - *Port Detailed Statistics* (**System > Statistics > System > Port Detailed**)
- Hover over any port to display status information for that port, as shown in the following illustration.



Example of Information Displayed by Hovering over a Port

Navigation Menu

The navigation menu, located at the top right of each UI page, lists the device's main features: **PoE**, **System**, **Switching**, **Routing**, **Security**, and **QoS**. You can access each feature's UI pages using a series of cascading menus.

To access an individual UI page, click the corresponding feature tab in the navigation menu to display a menu of subcategories. Select a subcategory and repeat this process until you see the desired page, and then select the page to display it in the main window.


For example, the following illustration shows how to access the *IPv6 Network Connectivity* page: first, select the main feature (**System** tab); then, the appropriate subcategory (**Connectivity**); and finally, the desired page (**IPv6**).



Navigation Menu View – System Tab Submenu

Each menu option (subcategory or page name) that you select is highlighted (the color changes to a lighter shade of gray). When you select a page, the navigation menus and submenus are again hidden, and the selected page appears in the main window.


In addition to the navigation menu, you can use the tabs at the top left of each page to quickly navigate among related pages. For example, from the *System Resource Configuration* page, simply click the *ARP Cache* or *Resource Status* tabs to display those pages without having to access the navigation menu, as shown in the following illustration.



Page Selection Tabs on System Resource Configuration Page

Configuration and Status Fields

The main area of the screen displays fields that you use to configure the switch and monitor its status. Configuration options allow you to input information using text input boxes, or make selections from drop-down boxes, radio buttons, and check boxes. Status fields display read-only information related to the switch and its configuration.



Example of Configuration and Status Fields

Command Buttons

Many UI pages also contain command buttons. These buttons, which typically appear at the bottom of a page but can also appear in the configuration and status field area, are labeled with either text or icons. The following table lists the common command buttons found throughout the UI pages.

Common Command Buttons

Button Text ^{1,2}	Icon	Function
Add		Adds a new entry to a table.
Clear	–	Removes all entries from a table, resets statistical counters to the default value, or clears all the statistics counters and resets all switch summary and detailed statistics to default values.
Delete		Removes the selected entry from the running configuration.
Download		Downloads data.
Edit		Changes an existing entry.
Generate		Generates a security certificate, key, etc.
Initialize		Resets the 802.1X state machine on the associated interface to the initialization state.
Logout	–	Ends the session.
Re-Authenticate		Forces the associated interface to restart the authentication process.
Refresh		Refreshes the page with the most current information, or refreshes the DHCP lease.
Remove	–	Deletes the selected entries.
Reset		Resets a field to its default value.
Submit	–	Sends the updated configuration to the switch. Configuration changes take effect immediately, but changes are not retained across a power cycle unless you save them to the system configuration file. IMPORTANT: To retain changes across a power cycle (reboot), you must save the configuration to non-volatile memory, by navigating to System > Configuration Storage > Save and clicking Save .
Upload		Uploads data.

¹ This is either the text label on a button, or the text that appears when hovering over a button labeled with an icon.

² Button names may include additional text, such as: **Add Vendor Option**, **Clear Entries**, **Remove Last Rule**, etc.

Table Sorting

All tables on UI pages can be sorted by columns. By default, the information in a table is sorted in ascending order, using the leftmost column as primary sort. To change the default sort order, click the heading above the column you want to sort the table by. Successive clicks on the heading toggle between ascending and descending order.

For example, the following illustration shows the *Event Log* page in its default sort order (sorted by *Log Index*). To sort the table entries (rows) by the *Event Time* field, simply click the *Event Time* heading.

Click to sort by Event Time

Log Index	Type	Filename	Line	Task ID	Code	Event Time
1	ERROR	cnfgr_tally.c	199	E2A61844	0000009B	0d-00-01:00
2	EVENT	Crashed!	0	E3C08B4	00000000	0d-00-01:25

Column Headings in Table

Table Filtering

This feature allows you to specify a filter that limits which rows are displayed in a table. This is useful to reduce the contents of a long table to a specific set of items or even one particular item. To use this feature, type a string of one or more characters into the *Filter* field at the upper-right corner of the table, as shown in the following illustration. If any field of a table row contains a match for the filter string, that row is displayed in the table. Matching is not case-sensitive.



Filtering the Contents of a Table

Help Page Access

The *Help* icon appears in the upper right corner of each UI page (see the illustration **“EdgeSwitch UI Page Layout” on page 11**). Click the *Help* icon to open a new page with information on the various fields and command buttons on the active page. Online help pages are context-sensitive – the help topic is specific to the active page.



Help Icon

User-Defined Fields

User-defined fields can contain 1-159 characters, unless otherwise noted on the configuration UI page. All characters may be used except for the following (unless specifically noted in the feature's Help page):

\ < / > * | ?

Using the Command-Line Interface

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

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

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

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

For more information about the CLI, see the *EdgeSwitch CLI Command Reference Guide*.

The *EdgeSwitch CLI Command Reference* lists each command available from the CLI by the command name and provides a brief description of the command. Each command reference also contains the following information:

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

Each `show` command in this document also includes a description of the information displayed by the command.

Chapter 2: Configuring Power over Ethernet

Use the PoE feature menu to display and configure the switch's Power over Ethernet (PoE) settings. The PoE tab contains links to the following features:

- **[“Configuring PoE” on page 17](#)**

Configuring PoE

This page displays information about the PoE settings on the switch's interfaces and allows you to configure those settings.

To access the *Power Over Ethernet* page, click **PoE > PoE Configuration** in the navigation menu.

Interface	PoE Mode	PoE Output	Current	Voltage
O/1	54V auto	Off	Off	Off
O/2	54V auto	Off	Off	Off
O/3	54V auto	Off	Off	Off
O/4	54V auto	Off	Off	Off
O/5	54V auto	2.65W	51.02mA	52.11V
O/6	54V auto	Off	Off	Off
O/7	54V auto	2.27W	43.45mA	52.24V
O/8	54V auto	Off	Off	Off
O/9	54V auto	0.71W	13.67mA	52.14V
O/10	54V auto	Off	Off	Off

Power Over Ethernet

Power Over Ethernet Fields

Field	Description
<i>Interface</i>	The interface (slot/port) for which PoE information is displayed.
<i>PoE Mode</i>	The PoE mode on the interface: • Off PoE is disabled for the interface. • 54V auto Standard PoE/PoE+ (802.3af/at) with auto-sensing is enabled for the interface. • 24V passive Passive 24V PoE output is enabled for the interface. Note: The 24V passive setting causes power to be output immediately with no auto-sensing. Before applying the 24V passive setting, ensure that the connected device can accept passive 24V PoE power.
The following fields apply only to interfaces whose PoE mode is set to <i>54V auto</i> :	
<i>PoE Output</i>	The interface's current PoE output power in W
<i>Current</i>	The interface's current output current in mA
<i>Voltage</i>	The interface's current output voltage in V

Use the buttons to perform the following tasks:

- To edit an interface's PoE settings, select the interface, click **Edit**, and make the changes as needed. Then, click **Submit** to apply the settings.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

Chapter 3: Configuring System Information

Use the features in the System feature menu to define the switch's relationship to its environment. The System folder contains links to the following features:


- [**"Viewing ARP Cache" on page 20**](#)
- [**"Viewing Inventory Information" on page 21**](#)
- [**"Viewing the Dual Image Status" on page 22**](#)
- [**"Viewing System Resources" on page 23**](#)
- [**"Defining General Device Information" on page 25**](#)
- [**"Basic Switch Configuration" on page 52**](#)
- [**"Managing Logs" on page 53**](#)
- [**"Configuring Email Alerts" on page 60**](#)
- [**"Viewing Device Port Information" on page 65**](#)
- [**"Defining SNMP Parameters" on page 72**](#)
- [**"Viewing System Statistics" on page 80**](#)
- [**"Using System Utilities" on page 91**](#)
- [**"Managing SNMP Traps" on page 101**](#)
- [**"Managing the DHCP Server" on page 103**](#)
- [**"Configuring Time Ranges" on page 110**](#)
- [**"Configuring DNS" on page 113**](#)
- [**"Configuring SNTP Settings" on page 116**](#)

Viewing ARP Cache

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

The ARP cache can support 1,024 entries, although this size is user-configurable to any value less than 1,024. When multiple network interfaces are supported by a device, as is typical of a router, either a single ARP cache is used for all interfaces, or a separate cache is maintained per interface. While the latter approach is useful when network addressing is not unique per interface, this is not the case for Ethernet MAC address assignment so a single ARP cache is employed.

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



ARP Cache

ARP Cache Fields

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

Use the buttons to perform the following tasks:

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

Viewing Inventory Information

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

To display the inventory information, click **System > Summary > Inventory** page in the menu.



System Inventory Information

System Inventory Information Fields


Field	Description
<i>System Description</i>	The product name of this switch.
<i>Machine Type</i>	The hardware platform of this switch.
<i>Machine Model</i>	The product model number.
<i>Serial Number</i>	The unique serial number used to identify this switch.
<i>Burned In MAC Address</i>	The burned-in universally administered MAC address of this switch.
<i>Software Version</i>	The <i>release.version.maintenance</i> number of the code currently running on the switch. For example, if the release is 1, the version is 2 and the maintenance number is 4, the format is "1.2.4."

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

Viewing the Dual Image Status

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

To display the *Dual Image Status* page, click **System > Firmware > Status** in the navigation menu.



Dual Image Status

Dual Image Status Fields

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

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

For information about how to update or change system images, see [“Using System Utilities” on page 91](#).

Viewing System Resources

System Resource Status

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

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

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



System Resource Status

System Resource Status Fields

Field	Description
<i>Memory Usage section:</i>	
<i>Free Memory</i>	Displays the available free memory on the switch.
<i>Alloc Memory</i>	Displays the allocated memory for the switch.
<i>Task ID</i>	Displays the ID of running tasks.
<i>Task Name</i>	Displays the name of the running tasks.
<i>CPU Utilization Report section:</i>	
<i>5 Seconds</i>	The percentage amount of CPU utilization consumed by the corresponding task in the last 5 seconds.
<i>60 Seconds</i>	The percentage amount of CPU utilization consumed by the corresponding task in the last 60 seconds.
<i>300 Seconds</i>	The percentage amount of CPU utilization consumed by the corresponding task in the last 300 seconds.

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

System Resource Configuration

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

System Resource Configuration

System Resource Configuration Fields

Field	Description
<i>Rising Threshold</i>	The CPU rising utilization threshold in percentage. A 0 (zero) percent threshold indicates that the CPU Utilization Notification feature is disabled.
<i>Rising Threshold Interval</i>	The CPU rising threshold interval in seconds. The time interval is configured in multiples of 5. A time interval of 0 (zero) seconds indicates that the CPU Utilization Notification feature is disabled.
<i>Falling Threshold</i>	The CPU falling utilization threshold in percentage. Configuration of this field is optional. If configured, the falling threshold value must be equal to or less than the rising threshold value. If not configured, it takes the same value as the rising threshold.
<i>Falling Threshold Interval</i>	The CPU falling threshold interval in seconds. Configuration of this field is optional. If configured, the falling interval value must be equal to or less than the rising interval value. If not configured, it takes the same value as the rising interval. The time interval is configured in multiples of 5.
<i>Free Memory Threshold</i>	The CPU free memory threshold in kilobytes. A 0 (zero) threshold value indicates that the CPU Free Memory Notification feature is disabled.

Use the buttons to perform the following tasks:

- Click **Submit** to apply the settings immediately to the running configuration.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

Defining General Device Information

The System menu's Configuration and Summary submenus contains links to pages that allow you to configure device parameters, such as the following features:

- [**"System Description" on page 26**](#)
- [**"IP Address Conflict Detection" on page 27**](#)
- [**"Network Connectivity" on page 27**](#)
- [**"Network Port IPv6 Neighbors" on page 30**](#)
- [**"DHCP Client Options" on page 31**](#)
- [**"HTTP Configuration" on page 31**](#)
- [**"Secure HTTP Configuration" on page 32**](#)
- [**"SSH Configuration" on page 33**](#)
- [**"Telnet Session Configuration" on page 34**](#)
- [**"User Accounts" on page 35**](#)
- [**"Authentication Server Users" on page 37**](#)
- [**"Logged in Sessions" on page 39**](#)
- [**"User Domain Name" on page 39**](#)
- [**"Accounting List" on page 40**](#)
- [**"Accounting Selection" on page 41**](#)
- [**"Authentication List Configuration" on page 42**](#)
- [**"Authentication Selection" on page 44**](#)
- [**"Line Password Configuration" on page 44**](#)
- [**"Enable Password Configuration" on page 45**](#)
- [**"Password Rules" on page 46**](#)
- [**"Last Password Result" on page 48**](#)
- [**"Denial of Service Configuration" on page 49**](#)
- [**"CLI Banner Configuration" on page 51**](#)

System Description

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

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

The screenshot shows the 'System Description' page with the following fields and values:

Field	Value
System Description	EdgeSwitch 24-Port 250W, v1.0.0.40B4592, Linux 3.6.5-14a26ed5
System Name	UBNT EdgeSwitch
System Location	
System Contact	
IP Address	10.0.2.145
System Up Time	0 days, 5 hours, 14 mins, 17 secs
Current SNTP Synchronized Time	Jul 18 21:34:21 2014 UTC

System Description

System Description Fields

Field	Description
<i>System Description</i>	The product name of this switch.
<i>System Name</i>	Enter the name you want to use to identify this switch. You may use up to 31 alphanumeric characters. This field is blank by default.
<i>System Location</i>	Enter the location of this switch. You may use up to 31 alphanumeric characters. This field is blank by default.
<i>System Contact</i>	Enter the contact person for this switch. You may use up to 31 alphanumeric characters. This field is blank by default.
<i>IP Address</i>	The IP Address assigned to the network interface. The network interface is the logical interface that allows remote management of the device via any of the front-panel switch ports. To change the IP address, see “Network Connectivity” on page 27 .
<i>System Up Time</i>	Displays the number of days, hours, and minutes since the last system restart.
<i>Current SNTP Synchronized Time</i>	Displays currently synchronized SNTP time in UTC. If no SNTP server has been configured and the time is not synchronized, this field displays “Not Synchronized.” To specify an SNTP server, see “Configuring SNTP Settings” on page 116 .

Defining System Information

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

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

Click **Refresh** to refresh the page with the most current data from the switch. Click **Cancel** to exit the page. To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

IP Address Conflict Detection

Use the *IP Address Conflict Detection* page to run the IP Address Conflict Detection tool, which detects IP address conflicts for IPv4 addresses. When a conflict is detected, the switch updates the status on the page, generates an SNMP trap, and logs a message noting the conflict.

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



IP Address Conflict Detection

IP Address Conflict Detection Fields

Field	Description
<i>IP Address Conflict Currently Exists</i>	Shows whether a conflicting IP address has been detected since status was last reset. False No conflict detected (the subsequent fields on this page are displayed as N/A). True Conflict was detected (the subsequent fields on this page show the relevant information).
<i>Last Conflicting IP Address</i>	The IP address of the interface that was last found to be in conflict. If multiple conflicts were detected, only the most recent occurrence is displayed. This field displays only if a conflict has been detected since the switch was last reset.
<i>Last Conflicting MAC Address</i>	The MAC address of the remote host associated with the IP address that was last found to be in conflict. If multiple conflicts are detected, only the most recent occurrence is displayed. This field is displayed only if a conflict has been detected since the switch was last reset.
<i>Time Since Conflict Detected</i>	The time elapsed (displayed in days, hours, minutes, and seconds) since the last address conflict was detected (provided Clear History has not yet been clicked). This field is displayed only if a conflict has been detected since the switch was last reset.

Use the buttons to perform the following tasks:

- To run the tool and check for possible address conflicts, click **Run Detection**.
- To reset the last IP address conflict detection status information seen by the switch, click **Clear History**.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Network Connectivity

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

The *IPv4 Network Connectivity* and *IPv6 Network Connectivity* pages allow you to change the IPv4 and IPv6 information using the EdgeSwitch UI. To access the pages, click **System > Connectivity > IPv4** or **IPv6** in the navigation menu.

The screenshot shows the 'IPv4 Network Connectivity' configuration window. At the top, there are tabs for 'IPv4', 'IPv6', 'IPv6 Neighbors', and 'DHCP Client Options'. The 'IPv4' tab is active. The configuration fields are as follows:

- Network Configuration Protocol:** Radio buttons for 'None', 'Bootp', and 'DHCP' (selected). A refresh icon is next to 'DHCP'.
- DHCP Client Identifier:** A checkbox that is currently unchecked.
- IP Address:** Text input field containing '10.0.2.145'.
- Subnet Mask:** Text input field containing '255.255.255.0'.
- Default Gateway:** Text input field containing '10.0.2.1'.
- MAC Address Type:** Radio buttons for 'Burned In' (selected) and 'Locally Administered'.
- Burned In MAC Address:** Text input field containing '04:18:D6:31:59:F4'.
- Locally Administered MAC Address:** Text input field containing '00-00-00-00-00-00'.
- Management VLAN ID:** Text input field containing '1'.

At the bottom of the form are three buttons: 'Submit', 'Refresh', and 'Cancel'. A copyright notice '© Copyright 2012-2014 Ubiquiti Networks, Inc.' is visible in the bottom right corner of the window.

IPv4 Network Connectivity

IPv4 Network Connectivity Fields

Field	Description
<i>Network Configuration Protocol</i>	Specifies what the switch should do following power-up. The factory default is <i>None</i> . Options are as follows: None Do not send any requests following power-up. Bootp Transmit a Bootp request. DHCP Transmit a DHCP request. Click to refresh the DHCP lease.
<i>DHCP Client Identifier</i>	The DHCP Client Identifier (Option 61) is used by DHCP clients to specify their unique identifier. DHCP servers use this value to index their database of address bindings. This value must be unique for all clients in an administrative domain. The Client Identifier string is displayed beside the check box if DHCP is enabled on the port on which the Client Identifier option is selected (the UI page must be refreshed after this change is made).
<i>IP Address</i>	The IP address of the network interface. The factory default value is <i>0.0.0.0</i> . Note: Each part of the IP address must start with a number other than zero. For example, IP addresses 001.100.192.6 and 192.001.10.3 are not valid.
<i>Subnet Mask</i>	The IP subnet mask for the interface. The factory default value is <i>0.0.0.0</i> .
<i>Default Gateway</i>	The default gateway for the IP interface. The factory default value is <i>0.0.0.0</i> .
<i>MAC Address Type</i>	Specifies whether to use the burned-in or the locally administered MAC address for in-band connectivity. The factory default is <i>Burned In</i> .
<i>Burned-In MAC Address</i>	This read-only field displays the MAC address that is burned-in to the network card at the factory. This MAC address is used for in-band connectivity if you choose not to configure a locally administered address.
<i>Locally Administered MAC Address</i>	Specifies a locally administered MAC address for in-band connectivity instead of using the burned-in universally administered MAC address. In addition to entering an address in this field, you must also set the MAC address type to locally administered. Enter the address as twelve hexadecimal digits (6 bytes) with a colon between each byte. Bit 1 of byte 0 must be set to a 1 and bit 0 to a 0; i.e., byte 0 must have a value between x'40' and x'7F'.
<i>Management VLAN ID</i>	Specifies the management VLAN ID of the switch. It may be configured to any value from 1 to 4093. The management VLAN is used for management of the switch. This field is configurable for administrative users and read-only for other users.

IPv6 Network Connectivity

IPv6 Network Connectivity Fields

Field	Description
<i>IPv6 Mode</i>	Enables or disables IPv6 mode.
<i>Network Configuration Protocol</i>	Specifies whether the device should attempt to acquire network information from a DHCPv6 server. The factory default is <i>None</i> , which disables the DHCPv6 client on the network interface.
<i>IPv6 Stateless Address AutoConfig Mode</i>	Sets the IPv6 stateless address autoconfiguration mode on the network interface. Enabled The network interface can acquire an IPv6 address through IPv6 Neighbor Discovery Protocol (NDP) and the use of Router Advertisement messages. Disabled The network interface will not use the native IPv6 address autoconfiguration features to acquire an IPv6 address.
<i>DHCPv6 Client DUID</i>	The client identifier used by DHCPv6 Client when sending messages to the DHCPv6 Server. Displayed only if <i>IPv6 Network Configuration Protocol</i> is set to <i>DHCP</i> .
<i>IPv6 Gateway</i>	The default gateway for the IPv6 network interface. Use the buttons to perform the following: - Click this button to change the field's setting. - Click this button to reset the field to the default value.
<i>Static IPv6 Addresses</i>	The configured static IPv6 addresses. Use the buttons to perform the following: - Click this button to add an IPv6 address by configuring the <i>New IPv6 Address</i> and <i>EUI Flag</i> fields in the <i>Add IPv6 Address</i> dialog box. - To remove an IPv6 address, select it and then click this button. To remove <i>all</i> IPv6 addresses, click this button in the heading row.
<i>New IPv6 Address</i>	Specifies the IPv6 address being added.
<i>EUI Flag</i>	Sets the EUI flag while configuring a new IPv6 address when selected. The default is option not selected.
<i>Dynamic IPv6 Addresses</i>	The configured dynamic IPv6 addresses.
<i>Default IPv6 Routers</i>	The default IPv6 Router address(es).

Use the buttons to perform the following tasks:

- If you change any of the network connectivity parameters, click **Submit** to apply the settings immediately to the running configuration.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

Network Port IPv6 Neighbors

When IPv6 is enabled on the service port, and a ping is initiated to a neighbor, the neighbor is added to the cache (if successful). The *Network Port IPv6 Neighbors* page displays data on these ports.

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



Network Port IPv6 Neighbors

Network Port IPv6 Neighbors Fields

Field	Description
<i>IPv6 Address</i>	The IPv6 address of a neighbor device that has been reachable on the local link through the network interface.
<i>MAC Address</i>	The MAC address of the neighboring device.
<i>Type</i>	The type of the neighbor entry, which is one of the following: Static The neighbor entry is manually configured. Dynamic The neighbor entry is dynamically resolved. Local The neighbor entry is a local entry. Other The neighbor entry is an unknown entry.
<i>Is Router</i>	Indicates whether the neighbor is a router. The possible values are: True The neighbor device is a router. False The neighbor device is not a router.
<i>Neighbor State</i>	Specifies the state of the neighbor cache entry. Following are the states for dynamic entries in the IPv6 neighbor discovery cache: Reachable The neighbor is reachable through the network interface. Stale The neighbor is not known to be reachable, and the system will begin the process to reach the neighbor. Delay The neighbor is not known to be reachable, and upper-layer protocols are attempting to provide reachability information. Probe The neighbor is not known to be reachable, and the device is attempting to probe for this neighbor. Unknown The reachability status cannot be determined.
<i>Last Updated</i>	The amount of time that has passed since the neighbor entry was last updated.

Use the buttons to perform the following tasks:

- To add a network port static IPv6 neighbor entry, click **Add**, configure the settings, and click **Submit** to apply the changes.
- To remove entries, select each entry to remove, click **Remove**, and confirm the removal.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

DHCP Client Options

Use the *DHCP Client Options* page to configure DHCP client settings on the system. To access the *DHCP Client Options* page, click **System** > **Connectivity** > **DHCP Client Options** in the navigation menu.

DHCP Client Options

DHCP Client Options Fields

Field	Description
<i>DHCP Vendor Class ID Mode</i>	The VCI administrative mode (<i>Enable</i> or <i>Disable</i>). When enabled, the DHCP client includes the text configured as the <i>DHCP Vendor Class ID String</i> in DHCP requests.
<i>DHCP Vendor Class ID String</i>	The text string added to DHCP requests as Option-60; i.e., Vendor Class Identifier option.

Use the buttons to perform the following tasks:

- Click **Submit** to apply the settings immediately to the running configuration.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

HTTP Configuration

Use the *HTTP Configuration* page to configure the HTTP server settings on the system. To access the *HTTP Configuration* page, click **System** > **Management Access** > **HTTP** in the navigation menu.

HTTP Configuration

HTTP Configuration Fields

Field	Description
<i>HTTP Admin Mode</i>	Used to <i>Enable</i> (default) or <i>Disable</i> the HTTP administrative mode. If this field is set to <i>Disable</i> , access to the UI is limited to secure HTTP, which is disabled by default.
<i>HTTP Session Soft Timeout</i>	Specifies the inactivity timeout value for HTTP sessions, in the range of 1 to 60 minutes (0 corresponds to an infinite timeout). The default value is 5 minutes.

HTTP Configuration Fields (Continued)

Field	Description
<i>HTTP Session Hard Timeout</i>	Specifies the hard timeout value for HTTP sessions in the range of 1 to 168 hours (0 corresponds to an infinite timeout). The default is 24 hours. This timeout is unaffected by the activity level of the session.
<i>Maximum Number of HTTP Sessions</i>	Specifies the maximum allowable number of HTTP sessions, in the range of 0 to 16 sessions. The default value is 16.

Use the buttons to perform the following tasks:

- Click **Submit** to apply the settings immediately to the running configuration.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Secure HTTP Configuration

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

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

Secure HTTP Configuration

Secure HTTP Configuration Fields

Field	Description
<i>HTTPS Admin Mode</i>	Used to <i>Enable</i> or <i>Disable</i> the HTTPS administrative mode. When this mode is enabled, the device can be accessed through a web browser using the HTTPS protocol.
<i>TLS Version 1</i>	Used to <i>Enable</i> or <i>Disable</i> Transport Layer Security Version 1.0. When enabled, communication between the web browser on the administrative system and the web server on the device is sent through TLS 1.0.
<i>SSL Version 3</i>	Used to <i>Enable</i> or <i>Disable</i> Secure Sockets Layer Version 3.0. When enabled, communication between the administrative system's web browser and the device's web server is sent through SSL 3.0. SSL must be administratively disabled while downloading an SSL certificate file from a remote server to the device.
<i>HTTPS Port</i>	The TCP port number that HTTPS uses.
<i>HTTPS Session Soft Time Out (Minutes)</i>	The maximum time in minutes that a user logged into an HTTPS session can be inactive before being automatically logged out of the HTTPS session.
<i>HTTPS Session Hard Time Out (Hours)</i>	The maximum time in hours that a user connected to the device via an HTTPS session can be inactive before being automatically logged out, regardless of the amount of HTTPS activity that occurs.

Secure HTTP Configuration Fields (Continued)

Field	Description
Maximum Number of HTTPS Sessions	The maximum number of HTTPS sessions that can be connected to the device simultaneously.
Certificate Status	The status of the SSL certificate generation process. Present The certificate has been generated and is present on the device Absent Certificate is not available on the device Generation In Progress An SSL certificate is currently being generated. Use the buttons next to this field to perform the following:  Click this button to download an SSL certificate file from a remote system to the device. Note that to download SSL certificate files, SSL must be administratively disabled.  Click this button to generate an SSL certificate to use for secure communication between the web browser and the embedded web server on the device.  Click this button to delete an SSL certificate (button available only if an SSL certificate is present on the device).

Use the buttons to perform the following tasks:

- If you make changes to the page, click **Submit** to apply the changes to the running configuration.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

SSH Configuration

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

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



SSH Configuration

SSH Configuration Fields

Field	Description
SSH Admin Mode	Used to <i>Enable</i> or <i>Disable</i> the SSH server administrative mode. When this mode is enabled, the device can be accessed by using an SSH client on a remote system.
SSH Version 1	Select this option to enable the device's SSH server to accept connections from SSH clients using SSH-1 protocol. Clear this option to disable connections from clients using SSH-1 protocol.
SSH Version 2	Select this option to enable the device's SSH server to accept connections from SSH clients using SSH-2 protocol. Clear this option to disable connections from clients using SSH-2 protocol.
SSH Connections Currently in Use	The number of active SSH sessions between remote SSH clients and the SSH server on the device.
Maximum number of SSH Sessions Allowed	The maximum number of SSH sessions that may be connected to the device simultaneously.
SSH Session Timeout (minutes)	The SSH session inactivity timeout value. A connected user that does not exhibit any SSH activity for this amount of time is automatically disconnected from the device.
RSA Key Status	The status of the SSH-1 Rivest-Shamir-Adleman (RSA) key file or SSH-2 RSA key file (PEM Encoded) on the device, which might be <i>Present</i>, <i>Absent</i>, or <i>Generation in Progress</i>. Use the buttons as follows:  Click to download an SSH-1 RSA or SSH-2 RSA key file from a remote system. In the <i>Download Certificate</i> dialog box, select the file type to download, browse to the file location on the remote system, select the file, and click Begin Transfer. The <i>Status</i> field provides information about the file transfer.  Click to manually generate an RSA key on the device.  Click to delete an RSA key downloaded to the device or manually generated on the device.
DSA Key Status	The status of the SSH-2 Digital Signature Algorithm (DSA) key file (PEM Encoded) on the device, which might be <i>Present</i>, <i>Absent</i>, or <i>Generation in Progress</i>. Use the buttons as follows:  Click to download an SSH-2 DSA key file from a remote system. In the <i>Download Certificate</i> dialog box, select the file type to download, browse to the file location on the remote system, select the file, and click Begin Transfer. The <i>Status</i> field provides information about the file transfer.  Click to manually generate a DSA key on the device.  Click to delete a DSA key downloaded to the device or manually generated on the device.

Use the buttons to perform the following tasks:

- If you make changes to the page, click **Submit** to apply the changes to the running configuration.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

Telnet Session Configuration

Telnet is a terminal emulation TCP/IP protocol. ASCII terminals can be virtually connected to the local device through a TCP/IP protocol network. Telnet is an alternative to a local login terminal where a remote login is required. The switch supports up to five simultaneous Telnet sessions. All CLI commands can be used over a Telnet session.

The *Telnet Session Configuration* page lets you control inbound Telnet settings on the switch. Inbound Telnet sessions originate on a remote system and allow a user on that system to connect to the switch CLI. To display the *Telnet Session Configuration* page, click **System** > **Management Access** > **Telnet** in the navigation menu.

Telnet Session Configuration

Telnet Session Configuration Fields

Field	Description
<i>Admin Mode</i>	Used to <i>Enable</i> or <i>Disable</i> the Telnet administrative mode. When enabled, the device may be accessed through the Telnet port (23). Disabling this mode value disconnects all existing Telnet connections and shuts down the Telnet port in the device.
<i>Session Timeout (Minutes)</i>	Specifies how many minutes (from 1 to 160) a Telnet session can be inactive before it is logged off. The factory default is 5. Note: When you change the timeout value, it is immediately applied to all active and inactive sessions. Any sessions that have been idle longer than the new timeout value are disconnected immediately.
<i>Maximum Number of Sessions</i>	Specifies the maximum number of Telnet sessions that can be connected simultaneously. The maximum is 4, which is also the factory default.
<i>Allow New Sessions</i>	Select this option to permit new Telnet sessions until the maximum number allowed is reached. Clear this option to disable new Telnet sessions (but existing sessions are not disconnected).

Use the buttons to perform the following tasks:

- If you make changes to the page, click **Submit** to apply the changes to the running configuration.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

User Accounts

By default, the switch contains one user account with read/write privileges. This account's default user name is *ubnt* and its password is *ubnt*; both user name and password are case-sensitive.

If you log on to the switch with the default read/write account (*ubnt*), you can use the *User Accounts* page to assign passwords and set security parameters for that account. You can also add up to five additional accounts (either read-only or read/write). You can delete all accounts except for the default account.



Note: Only a user with read/write privileges may alter data on this screen.

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



User Accounts

User Accounts Fields

Field	Description
User Name	The unique ID or name identifying the user account.
Access Level	Indicates the access or privilege level for this user. The options are: Read Write The user can view and modify the configuration. Read Only The user can view the configuration but cannot modify any fields. Suspended The user exists but is not permitted to log on to the device.
Lockout Status	Displays a user's current lockout status (<i>True</i> or <i>False</i>). A user is locked out of the system after failing to supply the correct password within the maximum allowed number of logins defined by the <i>Lockout Attempts</i> field on the <i>Password Rules</i> page. A locked-out user cannot log in again until an administrator resets the account using the <i>Unlock User Account</i> field (see table "Add New User and Edit Existing User Dialog Box Fields" on page 36).
Password Override	Identifies the password override complexity status for this user. Enable The system does not check the strength of the password. Disable When configuring a password, it is checked against the Strength Check rules configured for passwords.
Password Expiration	Indicates the current expiration date (if any) of the password.

The *User Accounts* page also provides the capability to add, edit, and remove user accounts:

- To add a user, click **Add**. The *Add new user* dialog box opens; specify the new account information in the available fields, and click **Submit** to create the new account.
- To edit an existing user, select the user's check box or click the row to select the account and click **Edit**. The *Edit existing user* dialog box opens; modify the account information as needed, and click **Submit** to apply the changes.
- To remove one or more user accounts, select one or more table entries, click **Remove**, and click **OK** to delete the selected entries.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

The following table describes the fields in the *Add new user* and *Edit existing user* dialog boxes.

Add New User and Edit Existing User Dialog Box Fields

Field	Description
User Name	The unique name for the account. Configurable only from the <i>Add new user</i> dialog box. Valid user names can contain up to 32 alphanumeric characters, plus "-" (hyphen) and "_" (underscore), and are not case-sensitive.
Password	Enter the optional new or changed password for the account. The password characters are not displayed on the page, but are disguised in a browser-specific manner. Passwords must be from 8 to 64 characters in length, and are case-sensitive.
Confirm	Enter the password again, to confirm that you entered it correctly. The password characters are not displayed on the page, but are disguised in a browser-specific manner.

Add New User and Edit Existing User Dialog Box Fields (Continued)

Field	Description
<i>Access Level</i>	Indicates the access or privilege level for this user. The options are: Read Write The user can view and modify the configuration. Read Only The user can view the configuration but cannot modify any fields. Suspended The user exists but is not permitted to log on to the device.
<i>Lockout Status</i>	(Edit existing user dialog box only) Displays a user's current lockout status (<i>True</i> if user is locked out of the system after failing to log in successfully within the configured number of login attempts).
<i>Unlock User Account</i>	(Edit existing user dialog box only) Select this option to unlock a user account that has been locked out (<i>Lockout Status</i> is <i>True</i>).
<i>Password Override</i>	Identifies the password override complexity status for this user. Enable The system does not check the strength of the password. Disable When configuring a password, it is checked against the Strength Check rules configured for passwords.
<i>Password Strength</i>	Indicates the date when the user's password will expire. This is determined by the date the password was created and the number of days specified in the <i>Aging</i> setting on the <i>Password Rules</i> page.
<i>Encrypted password</i>	Select this option to encrypt the password before it is stored on the device.

Authentication Server Users

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



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

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

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

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

When 802.1X authentication is enabled on the ports and the authentication method is LOCAL, port access is allowed only to users in this database that provide the correct name and password.

To access the *Auth Server Users* page, click **System > Users > Auth Server Users** in the navigation menu.



Auth Server Users

The *Auth Server Users* page lists the users (*User Name* field) in the authentication server user database. The following table describes the fields in the *Add new user* and *Edit existing user* dialog boxes.

Add New User and Edit Existing User Fields

Field	Description
<i>User Name</i>	A unique name used to identify the user account. Configurable only from the <i>Add new user</i> dialog box.
<i>Password Required</i>	Select this option to indicate that the user must enter a password to be authenticated. If this option is cleared, the user is required only to enter a valid user name.
<i>Password</i>	Specify the password to associate with the user name (if required).
<i>Confirm</i>	Re-enter the password to confirm the entry.
<i>Encrypted</i>	Select this option to encrypt the password before it is stored on the device.

Use the buttons to perform the following tasks:

- To add a user to the local authentication server database, click **Add**, configure the settings, and click **Submit** to apply the changes.
- To change the password information for an existing user, select the user to update, click **Edit**, configure the settings, and click **Submit** to apply the changes.
- To delete a user from the database, select each user to delete, click **Remove**, and confirm the deletion.
- To remove all users from the database, click **Clear All Users**.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

When you add a new user or edit an existing user, a pop-up window opens to allow you to configure the user information. The following table describes the configurable fields in these pop-up windows.

Auth Server Users Configuration Fields

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

Logged in Sessions

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

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



Logged In Sessions

Logged In Sessions Fields

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

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

User Domain Name

Use this page to configure the domain name to send to the authentication server, along with the user name and password, to authenticate a user attempting to access the device management interface. Domain name authentication is supported when user authentication is performed by a RADIUS server or TACACS+ server.

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



User Domain Name

User Domain Name Fields

Field	Description
User Domain Name Mode	Used to <i>Enable</i> or <i>Disable</i> the administrative mode of domain name authentication on the device. When enabled, the domain name is included when the user name and password are sent to the authentication server. The domain name can be specified either by the user in the <i>User Name</i> field on the login screen in a domain-name\username format, or it can be specified by the <i>Domain Name</i> field.
Domain Name	The domain name sent to the authentication server if the user does not provide one in the <i>User Name</i> field during logon. When only the username is provided, the device sends the username as <i>domain-name\username</i> , where <i>domain-name</i> is the string configured in this field. Use the buttons as follows: - To configure the <i>Domain Name</i> field, click this button and specify the desired string. - To reset the field to the default value, click this button and confirm the action.

Use the buttons to perform the following tasks:


- If you make changes to the page, click **Submit** to apply the changes to the running configuration.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Accounting List

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

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



Accounting List

Accounting List Fields

Field	Description
Accounting Type	The type of accounting list, which is one of the following: • Commands Each CLI command executed by the user, along with the time the command was executed, is recorded and sent to an external AAA server. • EXEC User login and logout times are recorded and sent to an external AAA server.
List Name	The name of the accounting list. This field can be configured only when adding a new accounting list.
Record Type	Indicates when to record and send information about the user activity: • StartStop Accounting notifications are sent at the beginning and end of an exec session or user-executed command. User activity does not wait for the accounting notification to be recorded at the AAA server. • StopOnly Accounting notifications are sent at the end of an exec session or user-executed command.
Method Options	The method(s) used to record user activity. The possible methods are as follows: • TACACS+ Accounting notifications are sent to the configured TACACS+ server. • RADIUS Accounting notifications are sent to the configured RADIUS server.

Accounting List Fields (Continued)

Field	Description
List Type	The type of accounting list, which is one of the following: Default The list is preconfigured on the system. This type of list cannot be deleted, and only the <i>Method Options</i> and <i>Record Type</i> settings are configurable. Configured The list has been added by a user.
Access Line	The access method(s) that use the list for accounting user activity. The settings for this field are configured on the <i>Accounting Selection</i> page.
Accounting Methods – This section of the <i>Add New Accounting List</i> dialog box contains the fields that you use to configure the accounting methods for the accounting list.	
Available Methods	The accounting methods that can be used for the accounting list. Select the method in the <i>Available Methods</i> field and click  to move it to the <i>Selected Methods</i> field.
Selected Methods	The accounting methods currently configured for the list. If this field lists multiple methods, the methods are applied in the order listed – if the switch fails to send accounting notifications using the first method, it tries again using the second method, and so on. To remove a method from the list, select it and click  .

Use the buttons to perform the following tasks:

- To configure a new accounting list, click **Add**, configure the settings in the *Add New Accounting List* dialog box, and then click **Submit** to apply the new settings to the switch.
- To edit a list, select the list's entry, click **Edit**, configure the settings in the *Edit Accounting List* dialog box, and then click **Submit** to apply the settings to the switch. The available settings depend on the list type.
- To remove a non-default accounting list, click the entry's  button and confirm the action.
- To reset the *Method Options* for a default accounting list to the factory default values, click the entry's  button and confirm the action.


To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Accounting Selection

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

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

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



Accounting Selection

Accounting Selection Fields

Field	Description
<i>Terminal</i> – The access methods in this section are CLI-based.	
<i>Telnet</i>	The <i>Exec</i> accounting list and the <i>Commands</i> accounting list to apply to users who access the CLI using a Telnet session.
<i>SSH</i>	The <i>Exec</i> accounting list and the <i>Commands</i> accounting list to apply to users who access the CLI using a secure shell (SSH) session.
<i>Hypertext Transfer Protocol</i> – The access methods in this section are through a web browser.	
<i>HTTP</i>	The <i>Exec</i> accounting list and the <i>Commands</i> accounting list to apply to users who access the web-based management interface using HTTP.
<i>HTTPS</i>	The <i>Exec</i> accounting list and the <i>Commands</i> accounting list to apply to users who access the web-based management interface using secure HTTP.

Use the buttons to perform the following tasks:


- If you make changes to the page, click **Submit** to apply the changes to the running configuration.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Authentication List Configuration

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

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



Authentication List Configuration

The following table shows the fields for the *Authentication List Configuration* page.

Authentication List Configuration Fields

Field	Description
<i>List Name</i>	The name of the authentication list. This field can be configured only when adding a new authentication list.
<i>Access Type</i>	How the user accesses the system. This field can be configured only when a new authentication list is added, and only the Login and Enable access types can be selected. The access types are as follows: Login User EXEC-level management access to the command-line interface (CLI) using a Telnet or SSH session. Access at this level has a limited number of CLI commands available to view or configure the system. Enable Privileged EXEC-level management access to the CLI using a Telnet or SSH session. In Privileged EXEC mode, read-write users have access to all CLI commands. HTTP Management-level access to the web-based user interface using HTTP. HTTPS Management-level access to the web-based user interface using secure HTTP. Dot1x Port-based access to the network through a switch port that is controlled by IEEE 802.1X.
<i>Method Options</i>	The method(s) used to authenticate a user who attempts to access the management interface or network. The possible methods are as follows: Enable Uses the locally configured Enable password to verify the user's credentials. Local Uses the ID and password in the Local User database to verify the user's credentials. RADIUS Sends the user's ID and password to the configured RADIUS server to verify the user's credentials. TACACS+ Sends the user's ID and password to the configured TACACS+ server to verify the user's credentials. None No authentication is used. IAS Uses the local Internal Authentication Server (IAS) database for 802.1X port-based authentication.
<i>List Type</i>	The type of list, which is one of the following: Default The list is preconfigured on the system. This type of list cannot be deleted, and only the <i>Method Options</i> are configurable. Configured The list has been added by a user.
<i>Access Line</i>	The access method(s) that use the list for authentication. The settings for this field are configured on the <i>Authentication Selection</i> page.
<i>Authentication Methods</i> – This section of the <i>Add New Authentication List</i> dialog box contains the fields that you use to configure the authentication methods for the authentication list.	
<i>Available Methods</i>	The authentication methods that can be used for the authentication list. To set the authentication method, select the method from the <i>Available Methods</i> field and click  to move it to the <i>Selected Methods</i> field.
<i>Selected Methods</i>	The authentication methods currently configured for the list. If this field lists multiple methods, the methods are applied in the order listed – if user authentication fails using the first method, the device tries again using the second method, and so on. If the current method is <i>None</i> , no authentication is performed (user is granted unconditional access); therefore, <i>None</i> must be the last method in the list. To remove a method from the list, select it and click  to return it to the <i>Available Methods</i> field.

Use the buttons to perform the following tasks:

- To configure a new authentication list, click **Add**, configure the settings in the *Add New Authentication List* dialog box, and click **Submit** to apply the settings to the switch.
- To edit a list, select the list's entry, click **Edit**, configure the settings in the *Edit Authentication List* dialog box, and click **Submit** to apply the settings to the switch. Available settings depend on the list type.
- To remove a non-default authentication list, click the entry's  button and confirm the action.
- To reset the *Method Options* for a default authentication list to the factory default values, click the entry's  button and confirm the action.
- Click **Refresh** to update the information on the screen.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

To create a new authentication list, see **"Authentication Server Users" on page 37**. To assign users to a specific authentication list, see **"User Accounts" on page 35**. To configure the 802.1X port security users, see **"RADIUS Settings" on page 218**.

Authentication Selection

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

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

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

Authentication Selection

The following table shows the fields for the *Authentication Selection* page.

Authentication Selection Fields

Field	Description
<i>Console</i>	The <i>Login</i> authentication list and the <i>Enable</i> authentication list to apply to users who attempt to access the CLI using a connection to the console port.
<i>Telnet</i>	The <i>Login</i> authentication list and the <i>Enable</i> authentication list to apply to users who attempt to access the CLI using a Telnet session.
<i>SSH</i>	The <i>Login</i> authentication list and the <i>Enable</i> authentication list to apply to users who attempt to access the CLI using a secure shell (SSH) session.

Use the command buttons to perform the following tasks:

- Click **Submit** to update the switch with the values on the screen.
- Click **Refresh** to update the information on the screen.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Line Password Configuration

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

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



The screenshot shows the 'Line Password Configuration' web page. At the top, there are five tabs: 'Line Password' (selected), 'Enable Password', 'Password Rules', 'Last Password', and 'Reset Passwords'. Below the tabs, the title 'Line Password Configuration' is displayed. A note states: 'Passwords must be from 8 to 64 characters in length.' The main section is titled 'Line Mode' and 'Password / Confirm Password'. It contains three rows for 'Console', 'Telnet', and 'SSH', each with a checkbox and two password input fields. At the bottom, there are 'Submit', 'Refresh', and 'Cancel' buttons. A copyright notice '© Copyright 2013-2014 Ubiquiti Networks, Inc.' is visible in the footer.

Line Password Configuration

Line Password Configuration Fields

Field	Description
Line Mode	Any or all of the following passwords may be changed on this page by checking the adjacent box: • Console • Telnet • SSH
Password (8-64 characters)	Enter the new password for the corresponding <i>Line Mode</i> in this field. Be sure the password conforms to the allowed number of characters. The password characters are not displayed on the page, but are disguised in a browser-specific manner.
Confirm Password (8-64 characters)	Re-enter the new password for the corresponding <i>Line Mode</i> in this field. This must be the same value entered in the <i>Line Password</i> field. Be sure the password conforms to the allowed number of characters. The password characters are not displayed on the page, but are disguised in a browser-specific manner.

Use the buttons to perform the following tasks:


- If you make changes to the page, click **Submit** to apply the changes to the running configuration.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

Enable Password Configuration

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

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



The screenshot shows the 'Enable Password Configuration' web page. At the top, there are five tabs: 'Line Password', 'Enable Password' (selected), 'Password Rules', 'Last Password', and 'Reset Passwords'. Below the tabs, the title 'Enable Password Configuration' is displayed. The main section contains two input fields: 'Enable Password' and 'Confirm Enable Password', both with a placeholder '(8 to 64 characters)'. At the bottom, there are 'Submit', 'Refresh', and 'Cancel' buttons. A copyright notice '© Copyright 2013-2014 Ubiquiti Networks, Inc.' is visible in the footer.

Enable Password Configuration

Enable Password Configuration Fields

Field	Description
<i>Enable Password</i>	Specify the password all users must enter after executing the enable command at the CLI prompt. The password characters are not displayed on the page, but are disguised in a browser-specific manner.
<i>Confirm Enable Password</i>	Enter the password again to confirm it. The password characters are not displayed on the page, but are disguised in a browser-specific manner.

Use the buttons to perform the following tasks:

- If you make changes to the page, click **Submit** to apply the changes to the running configuration.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Password Rules

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

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

Line Password **Enable Password** **Password Rules** **Last Password** **Reset Passwords**

Password Rules

Minimum Length: (0 to 64)

Aging (Days): (1 to 365, 0 = Default, 0 = Disable)

History: (0 to 100)

Lockout Attempts: (0 to 5, 0 = Default, 0 = Disable)

Strength Check: ☒ Disable ☐ Enable

Minimum Number of Uppercase Letters: (0 to 16, 2 = Default, 0 = Disable)

Minimum Number of Lowercase Letters: (0 to 16, 2 = Default, 0 = Disable)

Minimum Number of Numeric Characters: (0 to 16, 2 = Default, 0 = Disable)

Minimum Number of Special Characters: (0 to 16, 2 = Default, 0 = Disable)

Maximum Number of Repeated Characters: (0 to 15, 0 = Default, 0 = Disable)

Maximum Number of Consecutive Characters: (0 to 15, 0 = Default, 0 = Disable)

Minimum Character Classes: (0 to 4, 4 = Default, 0 = Disable)

Exclude Keyword Name:

Table is Empty

© Copyright 2013-2014 Ubiquiti Networks, Inc.

Password Rules

Password Rules Fields

Field	Description
<i>Minimum Length</i>	Specifies the minimum number of characters required for a valid password.
<i>Aging (days)</i>	The number of days that a user password is valid from the time the password is set. Once a password expires, the user is required to enter a new password at the next login.
<i>History</i>	The number of previous passwords that are retained to prevent password reuse. This helps to ensure that a user does not attempt to reuse the same password too often.
<i>Lockout Attempts</i>	The number of local authentication attempts that are allowed to fail before the user account is automatically locked (the account remains locked until the lockout is reset by the administrator on the user account page).
<i>Strength Check</i>	Used to <i>Enable</i> or <i>Disable</i> the password strength checking feature. Enabling this feature forces the user to configure passwords that satisfy the strong password requirements defined by the following fields.
<i>Minimum Number of Uppercase Letters</i>	Specifies the minimum number of uppercase letters a password must include.
<i>Minimum Number of Lowercase Letters</i>	Specifies the minimum number of lowercase letters a password must include.
<i>Minimum Number of Numeric Characters</i>	Specifies the minimum number of numbers a password must include.
<i>Minimum Number of Special Characters</i>	The minimum number of special characters (non-alphanumeric, such as @, #, &) that a valid password must contain.
<i>Maximum Number of Repeated Characters</i>	The maximum number of characters of any type that can repeat in a valid password. Repetition means the same character occurring in succession anywhere in the password, such as 11, %%%, or EEEE.
<i>Maximum Number of Consecutive Characters</i>	Specifies the maximum number of characters belonging to a sequence that are allowed to occur in a valid password. Consecutive characters are defined as a sequential pattern of case-sensitive alphabetic or numeric characters, such as 2345, def, or YZ.
<i>Minimum Character Classes</i>	Specifies the minimum number of character classes that a valid password must contain. There are four character classes: uppercase, lowercase, numeric, and special characters. This field allows you to define strength checking criteria for all four classes, but require passwords to meet only some of them. The number of character classes that must be met is specified by this value.
<i>Exclude Keyword Name</i>	The list of keywords that a valid password must not contain. Excluded keyword checking is case-insensitive. Additionally, a password cannot contain the backwards version of an excluded keyword. For example, if <i>pass</i> is an excluded keyword, passwords such as <i>23passA2c</i>, <i>ssapword</i>, and <i>PASsword</i> are prohibited. Use the buttons to perform the following tasks:  Click this button to add a keyword to the list. Type the word to exclude in the <i>Exclude Keyword Name</i> field, and click Submit.  Click this button next to a keyword to remove the keyword from the list, and confirm the action. To remove <i>all</i> keywords from the list, click the button in the header row and confirm the action.

Use the buttons to perform the following tasks:

- If you make changes to the page, click **Submit** to apply the changes to the running configuration.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Last Password Result

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

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



Last Password Result

Last Password Result Fields

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

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

Denial of Service Configuration

Use the *Denial of Service Configuration* page to configure Denial of Service (DoS) control. The EdgeSwitch software provides support for classifying and blocking specific types of DoS attacks. You can configure your system to monitor and block the types of attacks listed in the table below.

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

Denial of Service Configuration

Denial of Service Configuration Fields

Field	Description
<i>TCP Settings</i> – These options help prevent the device and the network from attacks that exploit the TCP header size or the information in the TCP or UDP headers of packets that the device receives.	
<i>First Fragment</i>	When selected, this option allows the device to drop packets that have a TCP header smaller than the value configured in the Min TCP Hdr Size field.
<i>TCP Port</i>	When selected, this option allows the device to drop packets that have the TCP source port equal to the TCP destination port.
<i>UDP Port</i>	When selected, this option allows the device to drop packets that have the UDP source port equal to the UDP destination port.
<i>SIP=DIP</i>	When selected, this option allows the device to drop packets that have a source IP address equal to the destination IP address.

Denial of Service Configuration Fields (Continued)

Field	Description
<i>SMAC=DMAC</i>	When selected, this option allows the device to drop packets that have a source MAC address equal to the destination MAC address.
<i>TCP FIN and URG and PSH</i>	When selected, this option allows the device to drop packets that have TCP Flags FIN, URG, and PSH set and a TCP Sequence Number equal to 0.
<i>TCP Flag and Sequence</i>	When selected, this option allows the device to drop packets that have TCP control flags set to 0 and the TCP sequence number set to 0.
<i>TCP SYN</i>	When selected, this option allows the device to drop packets that have TCP Flags SYN set.
<i>TCP SYN and FIN</i>	When selected, this option allows the device to drop packets that have TCP Flags SYN and FIN set.
<i>TCP Fragment</i>	When selected, this option allows the device to drop packets that have a TCP payload where the IP payload length minus the IP header size is less than the minimum allowed TCP header size.
<i>TCP Offset</i>	When selected, this option allows the device to drop packets that have a TCP header Offset set to 1.
<i>Min TCP Hdr Size</i>	The minimum TCP header size allowed. If First Fragment DoS prevention is enabled, the device will drop packets that have a TCP header smaller than this configured value.
<i>ICMP Settings</i> – These options help prevent the device and the network from attacks that involve issues with the ICMP echo request packets (pings) that the device receives.	
<i>ICMP</i>	Enable this option to allow the device to drop ICMP packets that have a type set to ECHO_REQ (ping) and a payload size greater than the ICMP payload size configured in the <i>Max ICMPv4 Size</i> field.
<i>Max ICMPv4 Size</i>	The maximum allowed ICMPv4 packet size. If ICMP DoS prevention is enabled, the device will drop ICMPv4 ping packets that have a size greater than this configured maximum ICMPv4 packet size.
<i>ICMPv6</i>	Enable this option to allow the device to drop ICMP packets that have a type set to ECHO_REQ (ping) and a payload size greater than the ICMP payload size configured in the <i>Max ICMPv6 Size</i> field.
<i>Max ICMPv6 Size</i>	The maximum allowed IPv6 ICMP packet size. If ICMP DoS prevention is enabled, the switch will drop IPv6 ICMP ping packets that have a size greater than this configured maximum ICMPv6 packet size.
<i>ICMP Fragment</i>	Enable this option to allow the device to drop fragmented ICMP packets.

Use the buttons to perform the following tasks:

- If you change any of the DoS settings, click **Submit** to apply the changes to the running configuration.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

CLI Banner Configuration

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

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

CLI Banner Configuration

CLI Banner Configuration Fields

Field	Description
<i>CLI Banner Message</i>	Text area for creating, viewing, or updating the CLI banner message. To create the CLI banner message, type the desired message in the text area. If you reach the end of the line, the text wraps to the next line. The line might not wrap at the same location in the CLI. To create a line break (carriage return) in the message, press Enter on the keyboard. The line break in the text area will be at the same location in the banner message when viewed through the CLI.

Use the buttons to perform the following tasks:

- If you make changes to the page, click **Submit** to apply the changes to the running configuration.
- Click **Clear** to clear the CLI banner message from the device. After you click **Clear**, you must confirm the action. You can also clear the CLI banner by deleting the text in the *CLI Banner Message* field and clicking **Submit**.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Basic Switch Configuration

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

Switch Configuration

Use the *Switch Configuration* page to set the amount of time to keep a learned MAC address entry in the forwarding database, or to enable or disable flow control mode on the switch.

The forwarding database contains both static entries that are never aged out, and dynamically learned entries that are removed if they are not updated within a specified time interval. The *Switch Configuration* page allows you to specify this time interval for learned entries.

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

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

Switch Configuration

Switch Configuration Fields

Field	Description
802.3x Flow Control Mode	Used to <i>Enable</i> or <i>Disable</i> 802.3x flow control on the switch: Disable The switch does not send PAUSE frames if the port buffers become full. Enable The switch can send PAUSE frames to a peer device if the port buffers become full.
MAC Address Aging Interval	Used to specify the number of seconds a dynamic address should remain in the MAC address table after it has been learned.



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

Use the buttons to perform the following tasks:

- If you make changes to the page, click **Submit** to apply the changes to the running configuration.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

Managing Logs

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

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

Log Configuration

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

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

Buffered Log Configuration

Admin Mode: ☐ Disable ☒ Enable

Behavior: ☒ Wrap ☐ Stop on Full

Command Logger Configuration

Admin Mode: ☒ Disable ☐ Enable

Console Log Configuration

Admin Mode: ☐ Disable ☒ Enable

Severity Filter:

Persistent Log Configuration

Admin Mode: ☒ Disable ☐ Enable

Severity Filter:

Syslog Configuration

Admin Mode: ☒ Disable ☐ Enable

Local UDP Port: (1 to 65535)

Submit Refresh Cancel

© Copyright 2013-2014 Ubiquiti Networks, Inc.

Log Configuration

Log Configuration Fields

Field	Description
<i>Buffered Log Configuration section:</i>	
<i>Admin Mode</i>	Used to <i>Enable</i> or <i>Disable</i> logging to the buffered (RAM) log file.
<i>Behavior</i>	Specify what the device should do when the buffered log is full. It can either overwrite the oldest messages (<i>Wrap</i>) or stop writing new messages to the buffer (<i>Stop on Full</i>).
<i>Command Logger Configuration section:</i>	
<i>Admin Mode</i>	Used to <i>Enable</i> or <i>Disable</i> logging of the command-line interface (CLI) commands issued on the device.
<i>Console Log Configuration section:</i>	
<i>Admin Mode</i>	Used to <i>Enable</i> or <i>Disable</i> logging to any serial device attached to the host.
<i>Severity Filter</i>	Select the severity of the messages to be logged. All messages at and above the selected threshold are logged to the console. The severity can be one of the following: • Emergency (0) The device is unusable. • Alert (1) Action must be taken immediately. • Critical (2) The device is experiencing primary system failures. • Error (3) The device is experiencing non-urgent failures. • Warning (4) The device is experiencing conditions that could lead to system errors if no action is taken. • Notice (5) The device is experiencing normal but significant conditions. • Info (6) The device is providing non-critical information. • Debug (7) The device is providing debug-level information.
<i>Persistent Log Configuration section:</i>	
<i>Admin Mode</i>	Used to <i>Enable</i> or <i>Disable</i> logging to the persistent log. These messages are not deleted when the device reboots.
<i>Severity Filter</i>	Select the severity of the messages to be logged. All messages at and above the selected threshold are logged to the console. See the previous severity filter description for more information about each severity level.
<i>Syslog Configuration section:</i>	
<i>Admin Mode</i>	Used to <i>Enable</i> or <i>Disable</i> logging to configured syslog hosts. When the syslog admin mode is disabled the device does not relay logs to syslog hosts, and no messages will be sent to any collector/relay. When the syslog admin mode is enabled, messages will be sent to configured collectors/relays using the values configured for each collector/relay.
<i>Local UDP Port</i>	The UDP port on the local host from which syslog messages are sent.

Use the buttons to perform the following tasks:


- If you make changes to the page, click **Submit** to apply the changes to the running configuration.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Buffered Log

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

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



Buffered Log

Buffered Log Fields

Field	Description
<i>Log Index</i>	The position of the entry within the buffered log file. The most recent log message has a <i>Log Index</i> value of 1.
<i>Log Time</i>	The time the entry was added to the log.
<i>Severity</i>	The severity level associated with the log entry. The severity can be one of the following: Emergency (0) The device is unusable. Alert (1) Action must be taken immediately. Critical (2) The device is experiencing primary system failures. Error (3) The device is experiencing non-urgent failures. Warning (4) The device is experiencing conditions that could lead to system errors if no action is taken. Notice (5) The device is experiencing normal but significant conditions. Info (6) The device is providing non-critical information. Debug (7) The device is providing debug-level information.
<i>Component</i>	The component that issued the log entry.
<i>Description</i>	The text description for the log entry.

Use the buttons to perform the following tasks:

- Click **Clear Log** to clear the buffered log messages and reset the counters. The buffered log will be repopulated with new entries as they occur on the system.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Event Log

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

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

Log Index	Type	Filename	Line	Task ID	Code	Event Time
1	ERROR	cfg_r_tally.c	199	02A51644	00000098	0d:00:01:00
2	EVENT	Crashed!	0	03C60804	00000000	0d:00:01:25
3	EVENT	usmdb_sim.c	3612	047A995C	00000000	0d:00:02:18
4	EVENT	usmdb_sim.c	3612	0483FC6C	00000000	0d:00:15:02
5	ERROR	cfg_r_tally.c	199	02C39644	00000098	0d:00:01:00
6	EVENT	Crashed!	0	02D9880C	00000000	0d:00:01:00

Event Log

Event Log Fields

Field	Description
<i>Log Index</i>	A display row index number used to identify the event log entry, with the most recent entry listed first (lowest number).
<i>Type</i>	The incident category that indicates the cause of the log entry: EVENT, ERROR, etc.
<i>Filename</i>	The EdgeSwitch source code filename identifying the code that detected the event.
<i>Line</i>	The line number within the source file of the code that detected the event.
<i>Task ID</i>	A system identifier of the task that was running when the event occurred. This value is assigned by, and is specific to, the operating system.
<i>Code</i>	An event-specific code value that is passed to the log handler by the source code file reporting the event.
<i>Event Time</i>	A time stamp (days, hours, minutes, and seconds) indicating when the event occurred, measured from the time the device was last reset. The only correlation between any two entries in the event log is the relative amount of time after a system reset that the event occurred.

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

Logging Hosts

Use the *Logging Hosts* page to configure remote logging hosts to which the switch can send logs.

To access the *Logging Hosts* page, click **System > Logs > Hosts** in the navigation menu. The *Logging Hosts* page is shown below.



Logging Hosts

Logging Hosts Fields

Field	A
<i>Host (IP Address/Host Name)</i>	The IP address or DNS-resolvable host name of the remote host to receive log messages. This field is not configurable when you click Edit .
<i>Status</i>	Indicates whether the host has been configured to be actively logging or not.
<i>Port</i>	The UDP port on the logging host to which syslog messages will be sent. The default port is 514. Specify the port in the text field.
<i>Severity Filter</i>	Use the menu to select the severity level threshold for log messages. Logs with a severity level at or above the configured level are forwarded to the host. For example, if you select Error, the logged messages include Error, Critical, Alert, and Emergency. The default severity level is Alert (1). The severity can be one of the following levels: Emergency (0) The highest warning level. If the device is down or not functioning properly, an emergency log is saved to the device. Alert (1) The second highest warning level. An alert log is saved if there is a serious device malfunction, such as all device features being down. Critical (2) The third highest warning level. A critical log is saved if a critical device malfunction occurs, for example, two device ports are not functioning, while the rest of the device ports remain functional. Error (3) A device error has occurred, such as if a port is offline. Warning (4) The lowest level of a device warning. Notice (5) Provides the network administrators with device information. Info (6) Provides device information. Debug (7) Provides detailed information about the log. Debugging should only be entered by qualified support personnel.

Use the buttons to perform the following tasks:

- To add a logging host, click **Add**. In the *Add Host* dialog box, fill in the *IP Address/Host Name*, *Port*, and *Severity Filter* fields, and click **Submit** to apply the changes to the switch's running configuration.
- To change information for an existing logging host, select the entry and click **Edit**. In the *Edit Host* dialog box, edit the *Port* and *Severity Filter* fields, and click **Submit** to apply the changes. You cannot edit the *IP Address/Host Name* field.
- To delete a configured logging host from the list, select the check box associated with each entry to delete, click **Remove**, and confirm the deletion.
- Click **Refresh** to update the screen and associated messages.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Syslog Source Interface Configuration

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

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

Syslog Source Interface Configuration

Syslog Source Interface Configuration Fields

Field	Description
<i>Type</i>	The type of interface to use as the source interface: None The primary IP address of the originating (outbound) interface is used as the source address. Interface The primary IP address of a physical port is used as the source address. VLAN The primary IP address of a VLAN routing interface is used as the source address. Tunnel The primary IP address of a tunnel interface is used as the source address.
<i>Interface</i>	When the selected <i>Type</i> is <i>Interface</i> , select the physical port to use as the source interface.
<i>VLAN ID</i>	When the selected <i>Type</i> is <i>VLAN</i> , select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.
<i>Tunnel ID</i>	When the selected <i>Type</i> is <i>Tunnel</i> , select the tunnel interface to use as the source interface.

Use the buttons to perform the following tasks:

- If you make changes to the page, click **Submit** to apply the changes to the running configuration.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Persistent Log

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

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



Persistent Log

Persistent Log Fields

Field	Description
<i>Log Index</i>	The position of the entry within the buffered log file. The most recent log message always has a <i>Log Index</i> value of 1.
<i>Log Time</i>	The time the entry was added to the log.
<i>Severity</i>	The severity level associated with the log entry. The severity can be one of the following: Emergency (0) The device is unusable. Alert (1) Action must be taken immediately. Critical (2) The device is experiencing primary system failures. Error (3) The device is experiencing non-urgent failures. Warning (4) The device is experiencing conditions that could lead to system errors if no action is taken. Notice (5) The device is experiencing normal but significant conditions. Info (6) The device is providing non-critical information. Debug (7) The device is providing debug-level information.
<i>Component</i>	The component that has issued the log entry.
<i>Description</i>	The text description for the log entry.

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

Configuring Email Alerts

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

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

Email Alert Global Configuration

Use the *Email Alert Global Configuration* page to configure the common settings for log messages emailed by the switch. To access the page, click **System > Advanced Configuration > Email Alerts > Global** in the navigation menu.

Email Alert Global Configuration

Email Alert Global Configuration Fields

Field	Description
<i>Admin Mode</i>	The administrative mode of the feature: Enable The device can send email alerts to the configured SMTP server. Disable The device will not send email alerts.
<i>From Address</i>	Specifies the email address of the sender (the switch).
<i>Log Duration</i>	This duration in minutes specifies how often to send the noncritical messages to the SMTP Server. For example, if set to 30 , the noncritical messages are sent every 30 minutes.
<i>Urgent Messages Severity</i>	Configures the urgent severity level(s) for log messages (urgent log messages are sent immediately). Select a severity level to define that level and all higher levels as urgent. Severity levels are (from highest to lowest severity): Emergency Indicates system is unusable (highest severity) Alert Indicates action must be taken immediately Critical Indicates critical conditions Error Indicates error conditions Warning Indicates warning conditions Notice Indicates normal but significant conditions Info Indicates informational messages Debug Indicates debug-level messages (lowest severity)
<i>Non Urgent Messages Severity</i>	Configures the nonurgent severity level(s) for log messages (nonurgent log messages are collected and sent in a digest form at the time interval specified by the <i>Log Duration</i> field). Select a severity level to define that level, and all levels up to but not including the lowest urgent level, as nonurgent. Messages below the severity level you specify are not sent via email. See the Urgent Messages Severity field description for information about the severity levels.
<i>Traps Severity</i>	Configures the severity level for trap log messages. See the Urgent Messages Severity field description for information about the severity levels.

Use the buttons to perform the following tasks:

- If you make changes to the page, click **Submit** to apply the changes to the running configuration.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

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

Email Alert Server Configuration

Use the *Email Alert Server Configuration* page to add, edit, and remove information about the network SMTP (mail) server that handles email alerts sent from the switch.

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



Email Alert Server Configuration

Email Alert Server Configuration Fields

Field	Description
<i>Address</i>	The IPv4/IPv6 address or host name of the SMTP server that handles email alerts that the device sends.
<i>Port</i>	The TCP port that email alerts are sent to on the SMTP server.
<i>Security</i>	The type of authentication to use with the mail server, which can be <i>TLSv1</i> (SMTP over SSL) or <i>None</i> (no authentication is required).
<i>User Name</i>	If the <i>Security</i> is <i>TLSv1</i> , this field specifies the user name required to access the mail server.
<i>Password</i>	If the <i>Security</i> is <i>TLSv1</i> , this field specifies the password associated with the configured user name for mail server access. When adding or editing the server, you must retype the password to confirm that it is entered correctly.

Use the buttons to perform the following tasks:


- To add an SMTP server, click **Add**, configure the desired settings, and click **Submit** to apply the changes.
- To edit the information for an existing SMTP server (except the *Host Name* or *IP Address* field), select the check box next to the entry and click **Edit**. When finished editing, click **Submit** to apply the changes.
- To delete a configured SMTP server from the list, select the check box next to the entry to delete and click **Remove**.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

Email Alert Statistics

Use the *Email Alert Statistics* page to view information about email alerts that the switch has sent or attempted to send. The statistics are cleared when the system is reset.

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



Email Alert Statistics

Email Alert Statistics Fields

Field	Description
<i>Number of Emails Sent</i>	The number of email alert messages successfully sent since the counters were cleared or the system was reset
<i>Number of Emails Failed</i>	The number of email alert messages that failed to be sent since the counters were cleared or the system was reset
<i>Time Since Last Email Sent</i>	The time in days, hours, minutes, and seconds that has passed since the last email alert message was successfully sent

Use the buttons to perform the following tasks:

- To reset the values on the page to zero, click **Clear Counters**.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Email Alert Subject Configuration

Use the *Email Alert Subject Configuration* page to configure the subject line of the urgent and nonurgent email alert messages sent from the switch. To access the page, click **System > Advanced Configuration > Email Alerts > Subject** in the navigation menu.

Email Alert Subject Configuration

Email Alert Subject Configuration Fields

Field	Description
<i>Message Type</i>	Select the message type for which you want to configure the subject line: <i>Urgent</i> or <i>Nonurgent</i> .
<i>Email Subject</i>	Specify the text to be displayed in the subject of the email alert message.
<i>Remove</i>	To reset the email alert subject to the default value, select the <i>Remove</i> option associated with the message type to reset, and click Delete .

Use the buttons to perform the following tasks:

- If you make changes to the page, click **Submit** to apply the changes to the running configuration.
- To remove a configured *Email Subject*, select the *Remove* check box associated with the entry, click **Delete**, and confirm the deletion.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Email Alert To Address Configuration

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

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



Email Alert To Address Configuration

Email Alert To Address Configuration Fields

Field	Description
<i>Message Type</i>	Select the type of message for which you want to specify a recipient address: <i>Urgent</i> or <i>Nonurgent</i> .
<i>To Address</i>	Specify the email address to which the selected type of messages are sent.

Use the buttons to perform the following tasks:

- To add an email address to the list of email alert message recipients, click **Add**, configure the desired settings, and click **Submit** to apply the changes.
- To remove configured email address entries from the list, select the *Remove* check box next to each entry to delete, click **Remove**, and confirm the deletion.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Viewing Device Port Information

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

Port Summary

Use the *Port Summary* page to view and configure the settings for all physical ports and Link Aggregation Groups (LAGs) on the switch. LAGs are also known as port channels.

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

Port Summary

Port Summary Fields

Field	Description
<i>Interface</i>	Identifies the port or LAG associated with the information in this row of the table.
<i>Interface Index</i>	The interface index object value assigned by the IF-MIB. This value is used to identify the interface when managing the device using SNMP.
<i>Type</i>	The interface type, which is one of the following: Normal The port is a normal port (it is not a LAG member or configured for port mirroring). Trunk Member The port is a member of a LAG. Mirrored The port is configured as a monitoring port and is the source port in a port mirroring session. For more information on port monitoring and probe ports, see “Mirroring” on page 69. Probe The port is configured as a monitoring port and is the destination port in a port mirroring session. For more information on port monitoring and probe ports, see “Mirroring” on page 69.
<i>Admin Mode</i>	The interface’s administrative mode: <i>Enabled</i> (default) or <i>Disabled</i> . If a port or LAG is administratively disabled, it cannot forward traffic.
<i>Physical Mode</i>	If the interface is not a LAG, this field displays the port’s configured speed and duplex mode: Auto The duplex mode and speed will be set by the auto-negotiation process. The port’s maximum capability (full duplex and 100 Mbps) will be advertised. <Speed> Half Duplex The port speeds available from the menu depend on the platform on which the EdgeSwitch software is running and which port you select. In half-duplex mode, the transmissions are one-way; that is, the port does not send and receive traffic at the same time. <Speed> Full Duplex The port speeds available from the menu depend on the platform on which the EdgeSwitch software is running and which port you select. In half-duplex mode, the transmissions are two-way. In other words, the port can send and receive traffic at the same time. If the interface is a LAG, this field displays LAG.
<i>Physical Status</i>	The port speed and duplex mode for physical interfaces. The physical status is not reported for LAGs. When a port is down, the physical status is unknown.

Port Summary Fields (Continued)

Field	Description
<i>STP Mode</i>	The Spanning Tree Protocol (STP) Administrative Mode associated with the port or LAG. STP is a Layer-2 protocol that provides a tree topology for switches on a bridged LAN. STP allows a network to have redundant paths without the risk of network loops, by providing a single path between end stations on a network. The possible values for STP mode are: • Enabled Spanning tree is enabled for this port. • Disabled Spanning tree is disabled for this port. For more information about STP, see “Configuring Spanning Tree Protocol” on page 161.
<i>LACP Mode</i>	The administrative mode of the Link Aggregation Control Protocol (LACP). The mode must be enabled in order for the port to participate in Link Aggregation. This field can have the following values: • Enabled The port uses LACP for dynamic LAG configuration. When LACP is enabled, the port sends and receives LACP Protocol Data Units (PDUs) with its link partner to confirm that the external switch is also configured for link aggregation. • Disabled The port supports static LAG configuration only. This mode might be used when the port is connected to a device that does not support LACP. When a port is added to a LAG as a static member, it neither transmits nor receives LACP PDUs.
<i>Link Status</i>	Indicates whether the Link is up or down. The link is the physical connection between the port or LAG and the interface on another device.
<i>Edit Port Configuration</i> dialog box – Click Edit to display this dialog box with configurable <i>Admin Mode</i>, <i>Physical Mode</i>, <i>STP Mode</i>, and <i>LACP Mode</i> fields, plus the following configurable fields:	
<i>Link Trap</i>	Indicates whether the port will send an SNMP trap when link status changes. • Enable (Default) The system sends a trap when the link status changes. • Disable The system does not send a trap when the link status changes.
<i>Maximum Frame Size</i>	The maximum Ethernet frame size the interface supports or is configured to support. The maximum frame size includes the Ethernet header, CRC, and payload.
<i>Broadcast Storm Recovery Level</i>	The broadcast storm control threshold for the port. If broadcast traffic on the Ethernet port exceeds this threshold, the system blocks (discards) the broadcast traffic. To configure this threshold (disabled by default), click Enable, enter a threshold value, and select the units for the threshold: • % The threshold value specifies a percentage of port speed from 0 to 100 (default: 5). • pps The threshold value is in packets per second.
<i>Multicast Storm Recovery Level</i>	The multicast storm control threshold for the port. If multicast traffic on the Ethernet port exceeds this threshold, the system blocks (discards) the multicast traffic. To configure this threshold (disabled by default), click Enable, enter a threshold value, and select the units for the threshold: • % The threshold value specifies a percentage of port speed from 0 to 100 (default: 5). • pps The threshold value is in packets per second.
<i>Unicast Storm Recovery Level</i>	The unicast storm control threshold for the port. If unicast traffic on the Ethernet port exceeds this threshold, the system blocks (discards) the unicast traffic. To configure this threshold (disabled by default), click Enable, enter a threshold value, and select the units for the threshold: • % The threshold value specifies a percentage of port speed from 0 to 100 (default: 5). • pps The threshold value is in packets per second.

Use the command buttons as follows:

- To edit a port's settings, select the port and click **Edit**. In the *Edit Port Configuration* dialog box, change the settings as needed, and click **Submit** to apply the changes.
- Click **Refresh** to redisplay the page with the latest information.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Port Description

Use the *Port Description* page to configure a human-readable description of the port.

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

Port Description

Display **10** rows Showing 1 to 10 of 32 entries Filter:

☐	Interface	Physical Address	PortList Bit Offset	Interface Index	Port Description
☐	0/1	04:18:D6:31:5A:40	1	1	
☐	0/2	04:18:D6:31:5A:40	2	2	
☐	0/3	04:18:D6:31:5A:40	3	3	
☐	0/4	04:18:D6:31:5A:40	4	4	
☐	0/5	04:18:D6:31:5A:40	5	5	
☐	0/6	04:18:D6:31:5A:40	6	6	
☐	0/7	04:18:D6:31:5A:40	7	7	
☐	0/8	04:18:D6:31:5A:40	8	8	
☐	0/9	04:18:D6:31:5A:40	9	9	
☐	0/10	04:18:D6:31:5A:40	10	10	

First Previous 1 2 3 4 Next Last Refresh Edit

© 2014-2018 Ubiquiti Networks, Inc.

Port Description

Port Description Fields

Field	Description
<i>Interface</i>	Select the interface for which data is to be displayed or configured.
<i>Physical Address</i>	The MAC address of the specified interface.
<i>PortList Bit Offset</i>	The bit offset value which corresponds to the port when the MIB object type PortList is used to manage the switch in SNMP.
<i>Interface Index</i>	The interface index object value assigned by the IF-MIB. This value is used to identify the interface when managing the device by using SNMP.
<i>Port Description</i>	The description, if any, associated with the interface to help identify it. By default, there is no associated description.

Use the command buttons as follows:

- To edit a port's description, select the port and click **Edit**. In the *Edit Port Description* dialog box, type the description in the *Port Description* field, and click **Submit** to apply the changes.
- Click **Refresh** to redisplay the page with the latest information.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Cable Test

The cable test feature enables you to determine the cable connection status on a selected port. You can also obtain an estimate of the length of the cable connected to the port, if the PHY on the ports supports this functionality.



Note: The cable test feature is supported only for copper cable. It is not supported for optical fiber cable.

To access the *Cable Test* feature, click **System > Port > Cable Test**.

The page displays additional fields when you click **Test Cable**. The fields that are displayed depend on the cable test results.

Cable Test

Cable Test Fields

Field	Description
<i>Interface</i>	Select the port (with connected cable) to be tested.
<i>Failure Location Distance</i>	The estimated distance from the end of the cable to the failure location. Note: This field displays a value only if the <i>Cable Status</i> is <i>Open</i> or <i>Short</i> ; otherwise, this field is blank.
<i>Cable Length</i>	The estimated length of the cable in meters. If the cable length cannot be determined, <i>Unknown</i> is displayed. This field shows the range between the shortest estimated length and the longest estimated length. Note: This field displays a value only when the <i>Cable Status</i> is <i>Normal</i> ; otherwise, this field is blank.
<i>Cable Status</i>	This field is displayed after you click Test Cable and test results are available. Values include. Normal The cable is working correctly. Open The cable is disconnected or there is a faulty connector. Open and Short There is an electrical short in the cable. Cable status test failed The cable status could not be determined. The cable may in fact be working.

Select a port from the *Interface* drop-down menu and click **Test Cable** to display its status.

If the port has an active link while the cable test is run, the link can go down for the duration of the test. The test may take several seconds to run. The command returns a cable length estimate if this feature is supported by the PHY for the current link speed.



Note: If the link is down and a cable is attached to a 10/100 Ethernet adapter, the displayed *Cable Status* may be *Open* or *Short* because some Ethernet adapters leave unused wire pairs unterminated or grounded.

Mirroring

Port mirroring selects the network traffic for analysis by a network analyzer. This is done for specific ports of the switch. As such, many switch ports are configured as source ports and one switch port is configured as a destination port. You have the ability to configure how traffic is mirrored on a source port. Packets that are received on the source port, that are transmitted on a port, or are both received and transmitted, can be mirrored to the destination port.

The packet that is copied to the destination port is in the same format as the original packet on the wire. This means that if the mirror is copying a received packet, the copied packet is VLAN tagged or untagged as it was received on the source port. If the mirror is copying a transmitted packet, the copied packet is VLAN tagged or untagged as it is being transmitted on the source port.

Use the *Multiple Port Mirroring* page to define port mirroring sessions. To access the *Multiple Port Mirroring* page, click **System** > **Port** > **Mirroring** in the navigation menu.

Multiple Port Mirroring

Multiple Port Mirroring Fields

Field	Description
<i>Session ID</i>	Specifies the monitoring session.
<i>Mode</i>	The administrative mode for the selected port mirroring session. If the mode is <i>Disabled</i> , the configured source is not mirroring traffic to the destination.
<i>Destination</i>	The interface that receives traffic from all configured source ports. To edit this field, click to open the <i>Destination Configuration</i> dialog box (see "" on page 70 for more information).
<i>IP ACL</i>	The IP access-list ID or name attached to the port mirroring session.
<i>MAC ACL</i>	The MAC access-list name attached to the port mirroring session.
<i>Source</i>	Possible values are: VLAN All the member ports of this VLAN are mirrored. Interface The port(s) configured to send traffic to the destination port. You can configure up to 5 source ports for each port mirroring session.
<i>Direction</i>	The direction of traffic on the source port(s) that is sent to the probe port. Possible values are: Tx and Rx Both ingress and egress traffic. Rx Ingress traffic only. Tx Egress traffic only.

Use the buttons to perform the following tasks:

- To configure the administrative mode for a port mirroring session or to select an ACL for flow-based mirroring, click **Configure Session**, configure the desired settings, and click **Submit** to apply the changes.
- To configure the destination as Remote VLAN or probe port, click  in the *Destination* field and click **Submit** to apply the change.
- To configure a source, click **Configure Source**. You can configure the source as Remote VLAN, VLAN, or Interface, specify one or more source ports for the mirroring session, or determine which traffic is mirrored (Tx, Rx, or both). Then, click **Submit** to apply the changes.
- To remove one or more source ports from the port mirroring session, select the check box associated with each source port to remove, click **Remove Source**, and confirm the removal.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Configuring a Port Mirroring Session



Note: A port will be removed from a VLAN or LAG when it becomes a destination mirror.

1. From the *Port Mirroring* page, click **Configure Session** to display the *Session Configuration* dialog box.
2. In the *Mode* field, select **Enable** to enable port mirroring.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Configuring Port Mirroring Source Ports

1. From the *Port Mirroring* page, click **Configure Source** to display the *Source Configuration* dialog box.
2. Configure the fields shown in the table below.

Multiple Port Mirroring – Source Configuration Fields

Field	Description
<i>Session ID</i>	Specifies the monitoring session.
<i>Type</i>	The type of interface to use as the source: • None The source is not configured. • Remote VLAN The VLAN configured as the RSPAN VLAN is the source. In an RSPAN configuration, the remote VLAN is the source on the destination device that has a physical port connected to the network traffic analyzer. • VLAN Traffic to and from a configured VLAN is mirrored; that is, all packets sent and received on all physical ports that are members of the VLAN are mirrored. • Interface Traffic is mirrored from one or more physical ports on the device.
<i>Remote VLAN</i>	The VLAN that is configured as the RSPAN VLAN.
<i>VLAN ID</i>	The VLAN to use as the source. Traffic from all physical ports that are members of this VLAN is mirrored. This field is available only when the selected <i>Type</i> is <i>VLAN</i> .
<i>Available Source Port(s)</i>	The physical port or ports to use as the source. Press and hold CTRL to select multiple ports. This field is available only when the selected <i>Type</i> is <i>Interface</i> .
<i>Direction</i>	Select the type traffic monitored on the source port, which can be one of the following: • Tx/Rx Monitors transmitted and received packets. • Rx Monitors received packets only. • Tx Monitors transmitted packets only.

3. Click **Add** to apply the changes to the system.

The new port mirroring session is enabled for the unit and port, and the device is updated. The source port appears in the *Source Port* list on the *Multiple Port Mirroring* page.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Configuring the Port Mirroring Destination

1. On the *Port Mirroring* page, click  next to the *Destination* field text box.

The *Destination Configuration* dialog box opens.

2. Configure the fields shown in the table below.

Multiple Port Mirroring – Add Source Ports Fields

Field	Description
<i>Type</i>	The type of interface to use as the destination: • None The destination is not configured. • Remote VLAN Traffic is mirrored to the VLAN on the system that is configured as the RSPAN VLAN. In an RSPAN configuration, the destination should be the Remote VLAN on any device that does not have a port connected to the network traffic analyzer. • Interface Traffic is mirrored to a physical port on the local device. The interface is the probe port that is connected to a network traffic analyzer.
<i>Remote VLAN</i>	The VLAN that is configured as the RSPAN VLAN.
<i>Port</i>	Click the drop-down box to select the port to which traffic is mirrored. If the <i>Type</i> is <i>Remote VLAN</i> , the selected port is a reflector port. The reflector port is a trunk port that carries the mirrored traffic towards the destination device. If the <i>Type</i> is <i>Interface</i> , the selected port is the probe port that is connected to a network traffic analyzer.

3. Click **Submit** to save the changes.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Removing Port Mirroring Sources

1. On the *Port Mirroring* page, select the source port to be removed.

2. Select one or more source ports to remove from the session.

Use the **CTRL** key to select multiple ports to remove.

3. Click **Remove Source**, and then click **OK** to confirm the operation.

The selected source ports are removed from the port mirroring session, and the device is updated.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Defining SNMP Parameters

Simple Network Management Protocol (SNMP) provides a method for managing network devices. The device supports SNMP version 1, SNMP version 2, and SNMP version 3.

SNMP v1 and v2

The SNMP agent maintains a list of variables, which are used to manage the device. The variables are defined in the Management Information Base (MIB). The MIB presents the variables controlled by the agent. The SNMP agent defines the MIB specification format, as well as the format used to access the information over the network. Access rights to the SNMP agent are controlled by access strings.

SNMP v3

SNMP v3 also applies access control and a new traps mechanism to SNMPv1 and SNMPv2 PDUs. In addition, the User Security Model (USM) is defined for SNMPv3 and includes:

- **Authentication:** Provides data integrity and data origin authentication.
- **Privacy:** Protects against disclosure of message content. Cipher-Block-Chaining (CBC) is used for encryption. Either authentication is enabled on an SNMP message, or both authentication and privacy are enabled on an SNMP message. However privacy cannot be enabled without authentication.
- **Timeliness:** Protects against message delay or message redundancy. The SNMP agent compares the incoming message to the message time information.
- **Key Management:** Defines key generation, key updates, and key use.

The device supports SNMP notification filters based on Object IDs (OID). OIDs are used by the system to manage device features. SNMP v3 supports the following features:

- Security
- Feature Access Control
- Traps


Authentication or Privacy Keys are modified in the SNMPv3 User Security Model (USM).

Use the *SNMP* page to define SNMP parameters. To display the *SNMP* page, click **System** > **SNMP** in the navigation menu.

SNMP Community Configuration

Access rights are managed by defining communities on the SNMPv1, 2 Community page. When the community names are changed, access rights are also changed. SNMP Communities are defined only for SNMP v1 and SNMP v2.

Use the *SNMP Community Configuration* page to enable SNMP and Authentication notifications. To display the page, click **System** > **Advanced Configuration** > **SNMP** > **Community** in the navigation menu.



SNMP Community Configuration

SNMP Community Configuration Fields

Field	Description
<i>Community Name</i>	Community name used in SNMPv1/v2 packets. This is configured in the client and identifies the access the user may connect with.
<i>Security Name</i>	Identifies the Security entry that associates Communities and Groups for a specific access type.
<i>Group Name</i>	Identifies the Group associated with this Community entry.
<i>IP Address</i>	Specifies the IP address that can connect with this community.

Use the buttons to perform the following tasks:

- To add a community, click **Add Community**, configure the settings, and click **Submit** to apply the change.
- To add a community group, click **Add Community Group**, configure the settings, and click **Submit** to apply the change.
- To delete a configured community from the list, select the check box next to its entry, click **Remove**, and confirm the deletion.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

SNMP v1/v2 Trap Receivers Configuration

Use the *SNMP v1/v2 Trap Receivers* configuration page to configure settings for each SNMPv1 or SNMPv2 management host that will receive notifications about traps generated by the device. The SNMP management host is also known as the SNMP trap receiver.

To access the page, click **System > Advanced Configuration > SNMP > Trap Receiver V1/V2** from the navigation menu.



SNMP v1/v2 Trap Receivers

SNMP v1/v2 Trap Receivers Fields

Field	Description
<i>Host IP Address</i>	The IP address of the SNMP management host that will receive traps generated by the device.
<i>Community Name</i>	The name of the SNMP community that includes the SNMP management host and the SNMP agent on the device.
<i>Notify Type</i>	The type of SNMP notification to send the SNMP management host: Inform An SNMP message that notifies the host when a certain event has occurred on the device. The message is acknowledged by the SNMP management host. This type of notification is not available for SNMPv1. Trap An SNMP message that notifies the host when a certain event has occurred on the device. The message is not acknowledged by the SNMP management host.
<i>SNMP Version</i>	The version of SNMP to use, which is either SNMPv1 or SNMPv2.
<i>Timeout Value</i>	The number of seconds to wait for an acknowledgment from the SNMP management host before resending an inform message.
<i>Retries</i>	The number of times to resend an inform message that is not acknowledged by the SNMP management host.
<i>Filter</i>	The name of the filter for the SNMP management host. The filter is configured by using the CLI and defines which MIB objects to include or exclude from the view. This field is optional.
<i>UDP Port</i>	The UDP port on the SNMP management host that will receive the SNMP notifications. If no value is specified when configuring a receiver, the default UDP port value is used.

Use the buttons to perform the following tasks:

- To add an SNMP trap receiver and configure its settings, click **Add**. In the *Add SNMP v1/v2 Host* dialog box, enter the required information, and then click **Submit** to apply the changes.
- To delete one or more SNMP trap receivers from the list, select each entry to delete, click **Remove**, and confirm the deletion.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

SNMP v3 Trap Receivers Configuration

Use the *Trap Receiver v3* page to configure settings for each SNMPv3 management host (also known as the SNMP trap receiver) that will receive notifications about traps generated by the device.

To access the *Trap Receiver v3* configuration page, click **System > Advanced Configuration > SNMP > Trap Receiver V3** from the navigation menu.



SNMP v3 Trap Receivers

SNMP v3 Trap Receivers Fields

Field	Description
<i>Host IP Address</i>	The IP address of the SNMP management host that will receive traps generated by the device.
<i>User Name</i>	The name of the SNMP user that is authorized to receive the SNMP notification.
<i>Notify Type</i>	The type of SNMP notification to send the SNMP management host: Inform An SNMP message notifying the host when a certain event occurs on the device. The SNMP management host acknowledge the message. This notification type is not available for SNMPv1. Trap An SNMP message that notifies the host when a certain event has occurred on the device. The message is not acknowledged by the SNMP management host.
<i>Security Level</i>	The security level associated with the SNMP user, which is one of the following: No Auth No Priv No authentication and no data encryption (no security). Auth No Priv Authentication, but no data encryption. With this security level, users send SNMP messages that use an MD5 key/password for authentication, but not a DES key/password for encryption. Auth Priv Authentication and data encryption. With this security level, users send an MD5 key/password for authentication and a DES key/password for encryption.
<i>Timeout Value</i>	The number of seconds to wait for an acknowledgment from the SNMP management host before resending an inform message.
<i>Retries</i>	The number of times to resend an inform message that is not acknowledged by the SNMP management host.
<i>Filter</i>	The name of the filter for the SNMP management host. The filter is configured by using the CLI and defines which MIB objects to include or exclude from the view. This field is optional.
<i>UDP Port</i>	The UDP port on the SNMP management host that will receive the SNMP notifications. If no value is specified when configuring a receiver, the default UDP port value is used.

Use the buttons to perform the following tasks:

- To add an SNMP trap receiver and configure its settings, click **Add**. In the *Add SNMP v3 Host* dialog box, enter the required information, and then click **Submit** to apply the changes.
- To delete one or more SNMP trap receivers from the list, select each entry to delete and click **Remove**, and confirm the deletion.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

SNMP Access Control Group

Use the *SNMP Access Control Group* page to configure SNMP access control groups. These SNMP groups allow network managers to assign different levels of authorization and access rights to specific device features and their attributes. The SNMP group can be referenced by the SNMP community to provide security and context for agents receiving requests and initiating traps as well as for management systems and their tasks. An SNMP agent will not respond to a request from a management system outside of its configured group, but an agent can be a member of multiple groups at the same time to allow communication with SNMP managers from different groups. Several default SNMP groups are preconfigured on the system.

To access the page, click **System > Advanced Configuration > SNMP > Access Control Group** in the navigation menu.

Group Name	Context Name	SNMP Version	Security Level	Read	Write	Notify
DefaultRead		SNMP V1	No Auth No Priv	Default		Default
DefaultRead		SNMP V2	No Auth No Priv	Default		Default
DefaultRead		SNMP V3	No Auth No Priv	Default		Default
DefaultRead		SNMP V3	Auth No Priv	Default		Default
DefaultRead		SNMP V3	Auth Priv	Default		Default
DefaultSuper		SNMP V1	No Auth No Priv	DefaultSuper	DefaultSuper	DefaultSuper
DefaultSuper		SNMP V2	No Auth No Priv	DefaultSuper	DefaultSuper	DefaultSuper
DefaultSuper		SNMP V3	No Auth No Priv	DefaultSuper	DefaultSuper	DefaultSuper
DefaultWrite		SNMP V1	No Auth No Priv	Default	Default	Default
DefaultWrite		SNMP V2	No Auth No Priv	Default	Default	Default

SNMP Access Control Group

SNMP Access Control Group Fields

Field	Description
Group Name	The name that identifies the SNMP group.
Context Name	The SNMP context associated with the SNMP group and its views. A user or a management application specifies the context name to get the performance information from the MIB objects associated with that context name. The Context EngineID identifies the SNMP entity that should process the request (the physical router), and the <i>Context Name</i> tells the agent in which context it should search for the objects requested by the user or the management application.
SNMP Version	The SNMP version associated with the group.
Security Level	The security level associated with the group, which is one of the following: No Auth No Priv No authentication and no data encryption (no security). This is the only Security Level available for SNMPv1 and SNMPv2 groups. Auth No Priv Authentication, but no data encryption. With this security level, users send SNMP messages that use an MD5 key/password for authentication, but no DES key/password for encryption. Auth Priv Authentication and data encryption. With this security level, users send an MD5 key/password for authentication and a DES key/password for encryption.

SNMP Access Control Group Fields (Continued)

Field	Description
<i>Read</i>	The level of read access rights for the group. The menu includes the available SNMP views. When adding a group, select the check box to allow the field to be configured, then select the desired view that restricts management access to viewing the contents of the agent.
<i>Write</i>	The level of write access rights for the group. The menu includes the available SNMP views. When adding a group, select the check box to allow the field to be configured, then select the desired view that permits management read-write access to the contents of the agent but not to the community.
<i>Notify</i>	The level of notify access rights for the group. The menu includes the available SNMP views. When adding a group, select the check box to allow the field to be configured, then select the desired view that permits sending SNMP traps or informs.

Use the buttons to perform the following tasks:

- To add an SNMP group, click **Add**. In the *Add new Access Control Group* dialog box, enter the settings, and then click **Submit** to apply the changes.
- To remove one or more SNMP groups, select each entry to delete, click **Remove**, and confirm the deletion.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

SNMP User Security Model

The *SNMP User Security Model* page provides the capability to configure the SNMP V3 user accounts.

To access the page, click **System** > **Advanced Configuration** > **SNMP** > **User Security Model** in the navigation menu.



SNMP User Security Model

SNMP User Security Model Fields

Field	Description
<i>Engine ID Type</i>	(Add New SNMP User dialog box only) • Local The engine ID is the switch's own engine ID (the engine ID cannot be changed). • Remote The engine ID is for a remote device and needs to be entered.
<i>Engine ID</i>	Each SNMPv3 agent has an engine ID that uniquely identifies the agent in the device. If given this entry will be used only for packets whose engine ID is this. This field takes a hexadecimal string in the form 0102030405.
<i>User Name</i>	Specifies the name of the SNMP user being added for the User-based Security Model (USM). Each user name must be unique within the SNMP agent user list. A user name cannot contain any leading or embedded blanks.

SNMP User Security Model Fields (Continued)

Field	Description
<i>Group Name</i>	An SNMP group is a group to which hosts running the SNMP service belong. A group name parameter is simply the name of that group by which SNMP communities are identified. The use of a group name provides some security and context for agents receiving requests and initiating traps and does the same for management systems and their tasks. An SNMP agent won't respond to a request from a management system outside its configured group, but an agent can be a member of multiple groups at the same time. This allows for communications with SNMP managers from different groups.
<i>Authentication Method</i>	Specifies the authentication protocol to be used on authenticated messages on behalf of the specified user. • SHA SHA protocol will be used. • MD5 MD5 protocol will be used. • None No authentication will be used for this user.
<i>Password</i>	Specifies the password used to generate the key to be used in authenticating messages on behalf of this user. This parameter must be specified if the <i>Authentication Method</i> parameter is not set to <i>None</i> .
<i>Privacy</i>	Specifies the privacy protocol to be used on encrypted messages on behalf of the specified user. This parameter is only valid if the <i>Authentication Method</i> parameter is not set to <i>None</i> . • DES DES protocol will be used. • None No privacy protocol will be used.
<i>Authentication Key</i>	Specifies the password used to generate the key to be used in encrypting messages to and from this user. This parameter must be specified if the <i>Privacy</i> parameter is not set to <i>None</i> .

Use the buttons to perform the following tasks:

- To add a user, click **Add**. The *Add New SNMP User* dialog box opens. Specify the new account information in the available fields, and then click **Submit** to apply the changes.
- To remove a user, select one or more table entries, click **Remove**, and confirm the deletion.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

SNMP Trap Source Interface Configuration

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

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

SNMP Trap Source Interface Configuration

SNMP Trap Source Interface Configuration Fields

Field	Description
<i>Type</i>	The type of interface to use as the source interface: • None The primary IP address of the originating (outbound) interface is used as the source address. • Interface The primary IP address of a physical port is used as the source address. • VLAN The primary IP address of a VLAN routing interface is used as the source address. • Tunnel The primary IP address of a tunnel interface is used as the source address.
<i>Interface</i>	When the selected <i>Type</i> is <i>Interface</i> , select the physical port to use as the source interface.
<i>VLAN ID</i>	When the selected <i>Type</i> is <i>VLAN</i> , select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.
<i>Tunnel ID</i>	When the selected <i>Type</i> is <i>Tunnel</i> , select the tunnel interface to use as the source interface

Use the buttons to perform the following tasks:

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

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

Viewing System Statistics

The pages in the Statistics folder contain a variety of information about the number and type of traffic transmitted from and received on the switch.

Switch Detailed Statistics

The *Switch Statistics* page shows detailed statistical information about the traffic the switch handles.

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

Statistics	Transmit	Receive
Octets Without Error	5508472	2240478
Packets Without Errors	7347	13322
Packets Discarded	0	0
Unicast Packets	5882	4703
Multicast Packets	1389	2007
Broadcast Packets	77	4615

Status	FDB Entries	VLANs
Current Usage	31	1
Peak Usage	36	1
Maximum Allowed	10384	255
Static Entries	1	1
Dynamic Entries	30	0
Total Entries Deleted	N/A	0

System	
Interface	65
Time Since Counters Last Cleared	0d 00:47:25

Refresh Clear Counters

© Copyright 2013-2014 Ubiquiti Networks, Inc.

Switch Statistics

Switch Statistics Fields

Field	Description
<i>Statistics section:</i>	
<i>Octets Without Error</i>	The total number of octets (bytes) of data successfully transmitted or received by the processor (excluding framing bits but including FCS octets).
<i>Packets Without Errors</i>	The total number of packets including unicast, broadcast, and multicast packets, successfully transmitted or received by the processor.
<i>Packets Discarded</i>	The number of outbound (Transmit column) or inbound (Receive column) packets which were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher-layer protocol. A possible reason for discarding a packet could be to free up buffer space.
<i>Unicast Packets</i>	The number of subnetwork-unicast packets delivered to or received from a higher-layer protocol.
<i>Multicast Packets</i>	The total number of packets transmitted or received by the device that were directed to a multicast address. Note that this number does not include packets directed to the broadcast address.
<i>Broadcast Packets</i>	The total number of packets transmitted or received by the device that were directed to the broadcast address. Note that this number does not include multicast packets.

Switch Statistics Fields (Continued)

Field	Description
<i>Status section:</i>	
<i>Current Usage</i>	In the <i>FDB Entries</i> column, the value shows the number of learned and static entries in the MAC address table. In the <i>VLANs</i> column, the value shows the total number of static and dynamic VLANs that currently exist in the VLAN database.
<i>Peak Usage</i>	The highest number of entries that have existed in the MAC address table or VLAN database since the most recent reboot.
<i>Maximum Allowed</i>	The maximum number of statically configured or dynamically learned entries allowed in the MAC address table or VLAN database.
<i>Static Entries</i>	The current number of entries in the MAC address table or VLAN database that an administrator has statically configured.
<i>Dynamic Entries</i>	The current number of entries in the MAC address table or VLAN database that have been dynamically learned by the device.
<i>Total Entries Deleted</i>	The number of VLANs that have been created and then deleted since the last reboot. This field does not apply to the MAC address table entries.
<i>System section:</i>	
<i>Interface</i>	The interface index object value of the interface table entry associated with the Processor of this switch. This value is used to identify the interface when managing the device by using SNMP.
<i>Time Since Counters Last Cleared</i>	The amount of time in days, hours, minutes, and seconds, that has passed since the statistics for this device were last reset.

Use the command buttons to perform the following actions:

- Click **Refresh** to refresh the data on the screen with the present state of the data in the switch.
- Click **Clear Counters** to clear all the statistics counters, resetting all switch summary and detailed statistics to default values. The discarded packets count cannot be cleared.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Port Summary

The *Port Summary Statistics* page shows statistical information about the packets received and transmitted by each port and LAG.

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

Interface	Rx Good	Rx Errors	Rx Bcast	Tx Good	Tx Errors	Tx Collisions
Q/1	11830	0	4856	7588	0	0
Q/2	0	0	0	0	0	0
Q/3	0	0	0	0	0	0
Q/4	0	0	0	0	0	0
Q/5	0	0	0	0	0	0
Q/6	0	0	0	0	0	0
Q/7	0	0	0	0	0	0
Q/8	0	0	0	0	0	0
Q/9	0	0	0	0	0	0
Q/10	0	0	0	0	0	0

Port Summary Statistics

Port Summary Statistics Fields

Field	Description
<i>Interface</i>	Identifies the port or LAG.
<i>Rx Good</i>	The total number of inbound packets received by the interface without errors.
<i>Rx Errors</i>	The number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol.
<i>Rx Bcast</i>	The total number of good packets received that were directed to the broadcast address. Note that this number does not include multicast packets.
<i>Tx Good</i>	The total number of outbound packets transmitted by the interface to its Ethernet segment without errors.
<i>Tx Errors</i>	The number of outbound packets that could not be transmitted because of errors.
<i>Tx Collisions</i>	The best estimate of the total number of collisions on this Ethernet segment.

Use the buttons to perform the following tasks:

- Click **Clear Counters** to clear all the statistics counters, resetting all summary and detailed statistics for this switch to default values. The discarded packets count cannot be cleared.
- Click **Clear All Counters** to clear counters for all switches in the stack.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Port Detailed Statistics

The *Port Detailed Statistics* page displays a variety of per-port traffic statistics. To access the *Port Detailed* page, click **System > Statistics > System > Port Detailed** in the navigation menu.

The following illustration shows the fields on the *Port Detailed Statistics* page.

Switch

Port Summary

Port Detailed

Network 094396

Port Detailed Statistics

Interface

G/1

Maximum Frame Size

1518

Packet Lengths Received and Transmitted

64 Octets

1791

65-127 Octets

7551

128-255 Octets

3180

256-511 Octets

1384

512-1023 Octets

1286

1024-1518 Octets

3431

1519-32767 Octets

0

32768-65535 Octets

0

65536-131071 Octets

0

Basic

Transmit

Receive

Unicast Packets

8180

4933

Multicast Packets

1463

2132

Broadcast Packets

81

4932

Total Packets Received

9724

35798

Packets > 1518 Octets

0

0

RSTx Frame Frames

0

0

RCS Errors

0

0

Protocol

Transmit

Receive

STP PDUs

0

0

RSTP PDUs

0

0

MSTP PDUs

1456

0

OSPF PDUs

0

0

OSPF PDUs

0

0

OSPF PDUs

0

0

Advanced - Transmits

Total Transmitted Packets Discarded

0

Single Collision Frames

0

Multiple Collision Frames

0

Excessive Collision Frames

0

Undersize Errors

0

OSPF Failed Registrations

0

OSPF Failed Registrations

0

Advanced - Receive

Total Packets Received Not Forwarded

0

Total Packets Received With MAC Errors

0

Overruns

0

Alignment Errors

0

Checksums Received

0

Fragmentary Received

0

Undersize Received

0

Unacceptable Frame Type

0

Time Since Counter Last Cleared

10:00:00.000

Refresh

Clear Counters

Clear All Counters

© Copyright 2003-2011 Ubiquiti Networks, Inc.

Port Detailed Statistics

Port Detailed Statistics Fields

Field	Description
<i>Interface</i>	Use the drop-down menu to select the interface for which data is to be displayed or configured. For non-stacking systems, this field is Slot/Port.
<i>Maximum Frame Size</i>	The maximum Ethernet frame size the interface supports or is configured to support. The maximum frame size includes the Ethernet header, CRC, and payload.
<i>Packet Lengths Received and Transmitted section:</i>	
<i>64 Octets</i>	The total number of packets (including bad packets) received or transmitted that were 64 octets in length (excluding framing bits but including FCS octets).
<i>65-127 Octets</i>	The total number of packets (including bad packets) received or transmitted that were between 65 and 127 octets in length inclusive (excluding framing bits but including FCS octets).
<i>128-255 Octets</i>	The total number of packets (including bad packets) received or transmitted that were between 128 and 255 octets in length inclusive (excluding framing bits but including FCS octets).
<i>256-511 Octets</i>	The total number of packets (including bad packets) received or transmitted that were between 256 and 511 octets in length inclusive (excluding framing bits but including FCS octets).
<i>512-1023 Octets</i>	The total number of packets (including bad packets) received or transmitted that were between 512 and 1023 octets in length inclusive (excluding framing bits but including FCS octets).
<i>1024-1518 Octets</i>	The total number of packets (including bad packets) received or transmitted that were between 1024 and 1518 octets in length inclusive (excluding framing bits but including FCS octets).
<i>1519-1522 Octets</i>	The total number of packets (including bad packets) received or transmitted that were between 1519 and 1522 octets in length inclusive (excluding framing bits but including FCS octets).
<i>1523-2047 Octets</i>	The total number of packets (including bad packets) received or transmitted that were between 1523 and 2047 octets in length inclusive (excluding framing bits but including FCS octets).
<i>2048-4095 Octets</i>	The total number of packets (including bad packets) received or transmitted that were between 2048 and 4095 octets in length inclusive (excluding framing bits but including FCS octets).
<i>4096-9216 Octets</i>	The total number of packets (including bad packets) received or transmitted that were between 4096 and 9216 octets in length inclusive (excluding framing bits but including FCS octets).
<i>Basic section:</i>	
<i>Unicast Packets</i>	The <i>Transmit</i> column shows the total number of packets that higher-level protocols requested be transmitted to a subnetwork unicast address, including those that were discarded or not sent. The <i>Receive</i> column shows the number of subnetwork unicast packets delivered to a higher-layer protocol.
<i>Multicast Packets</i>	The <i>Transmit</i> column shows the total number of packets that higher-level protocols requested be transmitted to a multicast address, including those that were discarded or not sent. The <i>Receive</i> column shows the number of multicast packets delivered to a higher-layer protocol.
<i>Broadcast Packets</i>	The <i>Transmit</i> column shows the total number of packets that higher-level protocols requested be transmitted to a broadcast address, including those that were discarded or not sent. The <i>Receive</i> column shows the number of broadcast packets delivered to a higher-layer protocol.
<i>Total Packets (Octets)</i>	The total number of octets of data (including those in bad packets) transmitted or received on the interface (excluding framing bits but including FCS octets). This object can be used as a reasonable estimate of Ethernet utilization. If greater precision is desired, the <i>etherStatsPkts</i> and <i>etherStatsOctets</i> objects should be sampled before and after a common interval.
<i>Packets > 1518 Octets</i>	The total number of packets transmitted or received by this interface that were longer than 1518 octets (excluding framing bits, but including FCS octets) and were otherwise well formed. This counter has a maximum increment rate of 815 counts per sec at 10 Mb/s.
<i>802.3x Pause Frames</i>	The number of MAC Control frames transmitted or received by this interface with an opcode indicating the PAUSE operation. This counter does not increment when the interface is operating in half-duplex mode.
<i>FCS Errors</i>	The total number of packets transmitted or received by this interface that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, but had a bad Frame Check Sequence (FCS) with an integral number of octets.
<i>Protocol section:</i>	
<i>STP BPDUs</i>	The number of Spanning Tree Protocol (STP) Bridge Protocol Data Units (BPDUs) transmitted or received by the interface.
<i>RSTP BPDUs</i>	The number of Rapid STP BPDUs transmitted or received by the interface.
<i>MSTP BPDUs</i>	The number of Multiple STP BPDUs transmitted or received by the interface.

Port Detailed Statistics Fields (Continued)

Field	Description
<i>GVRP PDUs</i>	The number of Generic Attribute Registration Protocol (GARP) VLAN Registration Protocol (GVRP) PDUs transmitted or received by the interface.
<i>GMRP PDUs</i>	The number of GARP Multicast Registration Protocol (GMRP) PDUs transmitted or received by the interface.
<i>EAPOL Frames</i>	The number of Extensible Authentication Protocol (EAP) over LAN (EAPOL) frames transmitted or received by the interface for IEEE 802.1X port-based network access control.
<i>Advanced - Transmit section:</i>	
<i>Total Transmit Packets Discarded</i>	The sum of single collision frames discarded, multiple collision frames discarded, and excessive frames discarded.
<i>Single Collision Frames</i>	A count of the number of successfully transmitted frames on a particular interface for which transmission is inhibited by exactly one collision.
<i>Multiple Collision Frames</i>	A count of the number of successfully transmitted frames on a particular interface for which transmission is inhibited by more than one collision.
<i>Excessive Collision Frames</i>	A count of frames for which transmission on a particular interface fails due to excessive collisions.
<i>Underrun Errors</i>	The total number of frames discarded because the transmit FIFO buffer became empty during frame transmission.
<i>GMRP Failed Registrations</i>	The number of times attempted GMRP registrations could not be completed.
<i>GVRP Failed Registrations</i>	The number of times attempted GVRP registrations could not be completed.
<i>Advanced - Receive section:</i>	
<i>Total Packets Received Not Forwarded</i>	The number of inbound packets which were chosen to be discarded to prevent them from being delivered to a higher-layer protocol, even though no errors had been detected. One possible reason for discarding such a packet is to free up buffer space.
<i>Total Packets Received With MAC Errors</i>	The total number of inbound packets that contained errors preventing them from being delivered to a higher-layer protocol.
<i>Overruns</i>	The total number of frames discarded as this port was overloaded with incoming packets, and could not keep up with the inflow.
<i>Alignment Errors</i>	The total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, but had a bad Frame Check Sequence (FCS) with a non-integral number of octets.
<i>Jabbers Received</i>	The total number of packets received that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error). Note that this definition of jabber is different than the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition where any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms.
<i>Fragments Received</i>	The total number of packets received that were less than 64 octets in length with ERROR CRC (excluding framing bits but including FCS octets).
<i>Undersize Received</i>	The total number of packets received that were less than 64 octets in length with GOOD CRC (excluding framing bits but including FCS octets).
<i>Unacceptable Frame Type</i>	The number of frames discarded from this interface due to being a frame type that the interface cannot accept.
<i>Time Since Counters Last Cleared</i>	The amount of time in days, hours, minutes, and seconds, that has passed since the statistics for this interface were last reset.

Use the buttons to perform the following tasks:

- Click **Clear Counters** to clear all the counters. This resets all statistics for this port to the default values.
- Click **Clear All Counters** to clear all the counters for all ports on the switch. The button resets all statistics for all ports to default values.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Network Port DHCPv6 Client Statistics

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

To access the page, click **System > Statistics > System > Network DHCPv6** in the navigation menu.

Network Port DHCPv6 Client Statistics	
Advertisement Packets Received	0
Reply Packets Received	0
Received Advertisement Packets Discarded	0
Received Reply Packets Discarded	0
Malformed Packets Received	0
Total Packets Received	0
Solicit Packets Transmitted	0
Request Packets Transmitted	0
Renew Packets Transmitted	0
Rebind Packets Transmitted	0
Release Packets Transmitted	0
Total Packets Transmitted	0

Refresh Clear Counters

© Copyright 2013-2014 Ubiquiti Networks, Inc.

Network Port DHCPv6 Client Statistics

Network Port DHCPv6 Client Statistics Fields

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

Network Port DHCPv6 Client Statistics Fields (Continued)

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

Use the buttons to perform the following tasks:


- Click **Clear Counters** to clear all the counters. This resets all statistics for this port to the default values.
- Click **Refresh** to refresh the data on the screen and display the most current statistics.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

Time-Based Group Statistics

Use the *Time-Based Group Statistics* page to define criteria for collecting time-based statistics for interface traffic. The time-based statistics can be useful for troubleshooting and diagnostics purposes. The statistics application uses the system clock for time-based reporting, so it is important to configure the system clock (manually or through SNTP) before using this feature.

To access the page, click **System** > **Statistics** > **Time Based** > **Group** in the navigation menu.



Time-Based Group Statistics

Time-Based Group Statistics Fields

Field	Description
<i>Group</i>	The type of traffic statistics to collect for the group, which is one of the following: • Received The number of packets received on the interfaces within the group. • Received Errors The number of packets received with errors on the interfaces within the group. • Transmitted The number of packets transmitted by the interfaces within the group. • Received Transmitted The number of packets received and transmitted by the interfaces within the group. • Port Utilization The percentage of total bandwidth used by the port within the specified time period. • Congestion The percentage of time within the specified time range that the ports experienced congestion.
<i>Time Range</i>	The name of the periodic or absolute time range to use for data collection. The time range is configured using the <i>Time Range Entry Summary</i> page (see “Time Range Entry Configuration” on page 111). The time range must be configured on the system before the time-based statistics can be collected.

Time-Based Group Statistics Fields (Continued)

Field	Description
<i>Reporting Methods</i>	The methods for reporting the collected statistics at the end of every configured time range interval. The available options are: • None The statistics are not reported to the console or an external server. They can be viewed only by using the EdgeSwitch UI or by issuing a CLI command. • Console The statistics are displayed on the console. • E-Mail The statistics are sent to an e-mail address. The SMTP server and e-mail address information is configured by using the appropriate Email Alerts pages. • Syslog The statistics are sent to a remote syslog server. The syslog server information is configured on the Logging Hosts page.
<i>Interfaces</i>	The interface or interfaces on which data is collected. To select multiple interfaces when adding a new group, press and hold CTRL, and then click each interface to include in the group.

Use the buttons to perform the following tasks:

- To add a set of time-based traffic group statistics to collect, click **Add**, configure the desired settings, and then click **Submit** to apply the changes.
- To delete one or more time-based statistics groups, select each entry to delete and click **Remove**.
- Click **Refresh** to refresh the data on the screen with the present state of the data in the switch.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

Time-Based Flow Statistics

Use this page to define criteria for collecting time-based statistics for specific traffic flows. The statistics include a per-interface hit count based on traffic that meets the match criteria configured in a rule for the interfaces included in the rule. The hit count statistics are collected only during the specified time range. The statistics application uses the system clock for time-based reporting. Configure the system clock (manually or through SNTP) before using the time-based statistics feature.

To access the *Time-Based Flow Statistics* page, click **System** > **Statistics** > **Time Based** > **Flow Based** in the navigation menu.



Time-Based Flow Statistics

Time-Based Flow Statistics Fields

Field	Description
<i>Reporting Methods</i>	The methods for reporting the collected statistics at the end of every configured interval: • None The statistics are not reported to the console or an external server. They can be viewed only by using the EdgeSwitch UI or by issuing a CLI command. • Console The statistics are displayed on the console. • E-Mail The statistics are sent to an e-mail address. The SMTP server and e-mail address information is configured by using the appropriate <i>Email Alerts</i> pages. • Syslog The statistics are sent to a remote syslog server. The syslog server information is configured on the <i>Logging Hosts</i> page.  To change the reporting methods for all flow-based statistics rules, click this button and select one or more methods.  Click this button to reset the field to the default value.
<i>Rule Id</i>	The number that identifies the flow-based statistics collection rule.
<i>Time Range</i>	The name of the periodic or absolute time range to use for data collection. The time range is configured using the <i>Time Range Entry Summary</i> page (see “Time Range Entry Configuration” on page 111). The time range must be configured on the system before the time-based statistics can be collected.
<i>Match Conditions</i>	The criteria that a packet must meet to match the rule.
<i>Interfaces</i>	The interface or interfaces on which the flow-based rule is applied. Only traffic on the specified interfaces is checked against the rule.
When you click Add , the <i>Time Based Flow Configuration</i> dialog box opens and allows you to configure a rule for traffic flow statistics. The match conditions are optional, but the rule must specify at least one match condition. The match conditions are as follows:	
<i>Match All</i>	Select this option to indicate that all traffic matches the rule and is counted in the statistics. This option is exclusive to all other match criteria, so if <i>Match All</i> is selected, no other match criteria can be configured.
<i>Source IP</i>	The source IP address to match in the IPv4 packet header.
<i>Destination IP</i>	The destination IP address to match in the IPv4 packet header.
<i>Source MAC</i>	The source MAC address to match in the ingress frame header.
<i>Destination MAC</i>	The destination MAC address to match in the ingress frame header.
<i>Source TCP Port</i>	The TCP source port to match in the TCP header.
<i>Destination TCP Port</i>	The TCP destination port to match in the TCP header.
<i>Source UDP Port</i>	The UDP source port to match in the UDP header.
<i>Destination UDP Port</i>	The UDP destination port to match in the UDP header.

Use the buttons to perform the following tasks:

- To add a rule and define criteria for flow-based statistics that are collected within a time range, click **Add**. In the *Time Based Flow Configuration* dialog box, configure the settings, and then click **Submit** to apply the changes.
- To delete one or more flow-based rules for time-based statistics, select each entry to delete and click **Remove**.
- Click **Refresh** to refresh the data on the screen with the present state of the data in the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Time-Based Statistics

Use this page to view time-based statistics collected for the configured traffic groups and flow-based rules.

To access the *Time-Based Statistics* page, click **System** > **Statistics** > **Time Based** > **Statistics** in the navigation menu.



Time-Based Statistics

Time-Based Statistics Fields

Field	Description
<i>ID</i>	The traffic group name or flow-based rule ID associated with the rest of the statistics in the row.
<i>Interface</i>	The interface on which the statistics were reported.
<i>Counter Id</i>	For traffic group statistics, this field identifies the type of traffic.
<i>Counter Value</i>	For traffic group statistics, this field shows the number of packets of the type identified by the <i>Counter Id</i> field that were reported on the interface during the time range.
<i>Port Utilization</i>	For a port utilization traffic group, this field reports the percentage of the total available bandwidth used on the interface during the time range.
<i>Hit Count</i>	For flow-based statistics, this field reports the number of packets that matched the flow-based rule criteria during the time range.

Click **Refresh** to refresh the data on the screen with the present state of the data in the switch.

Using System Utilities

The System Utilities feature menu contains links to UI pages that help you configure features that help you manage the switch.

System Reset

Use the *System Reset* page to reboot the system. To access the *System Reset* page, click **System > Utilities > System Reset** in the navigation menu.




System Reset

Click **Reset** to initiate the system reset. If you have not saved the changes that you submitted since the last system reset, the changes will not be applied to the system after the reset.

Ping

Use the *Ping* page to tell the switch to send a Ping request to a specified IP address. You can use this feature to check whether the switch can communicate with a particular network host.

To access the *Ping* page, click **System > Utilities > Ping** in the navigation menu.



Ping

Ping Fields

Field	Description
<i>Host Name or IP Address</i>	Enter the IP address or the host name of the station you want the switch to ping. The initial value is blank. This information is not retained across a power cycle.
<i>Count</i>	The number of ICMP echo request packets to send to the host.
<i>Interval</i>	The number of seconds to wait between sending ping packets.
<i>Size</i>	The size of the ping packet, in bytes. Changing the size allows you to troubleshoot connectivity issues with a variety of packet sizes, such as large or very large packets.
<i>Source</i>	The source IP address or interface to use when sending the echo request packets. If source is not required, select None as the <i>Source</i> option.
<i>IP Address</i>	The source IP address to use when sending the Echo requests packets. This field is enabled when the <i>Source</i> option is set to <i>IP Address</i> .
<i>Interface</i>	The interface to use when sending the Echo requests packets. This field is enabled when the <i>Source</i> option is set to <i>Interface</i> .
<i>Status</i>	Displays the results of the ping.
<i>Results</i>	The results of the ping test, which includes information about the reply (if any) received from the host.

Use the buttons to perform the following tasks:

- Click **Start** to initiate the ping test. The device sends the specified number of ping packets to the host.
- Click **Stop** to interrupt the current ping test.
- Click **Refresh** to refresh the data on the screen with the present state of the data in the switch.

Ping IPv6

Use the *Ping IPv6* page to tell the device to send one or more ping requests to a specified IPv6 host. You can use the ping request to check whether the device can communicate with a particular host on an IPv6 network. A ping request is an Internet Control Message Protocol version 6 (ICMPv6) echo request packet. The information you enter on this page is not saved as part of the device configuration.

To access the *Ping IPv6* page, click **System** > **Utilities** > **Ping IPv6** in the navigation menu.

Ping IPv6

Ping IPv6 Fields

Field	Description
<i>Ping</i>	Select either a <i>Global</i> IPv6 address or a <i>Link Local</i> address to ping. A global address is routable over the Internet, while a link-local address is intended for communication only within the local network. Link local addresses have a prefix of <i>fe80::/64</i> .
<i>Interface</i>	This field is displayed only when <i>Link Local</i> is selected. Select an IPv6 interface to initiate the ping.
<i>Host Name or IPv6 Address</i>	Enter the global or link-local IPv6 address, or the DNS-resolvable host name of the station to ping. If the ping type is <i>Link Local</i> , you must enter a link-local address and cannot enter a host name.
<i>Count</i>	Enter the number of ICMP echo request packets to send to the host.
<i>Interval</i>	Enter the number of seconds to wait between sending ping packets.
<i>Size</i>	The size of the ping packet, in bytes. Changing the size allows you to troubleshoot connectivity issues with a variety of packet sizes, such as large or very large packets.
<i>Source</i>	The source IP address or interface to use when sending the echo request packets. If source is not required, select None as the <i>Source</i> option.
<i>IPv6 Address</i>	The source IPv6 address to use when sending the Echo requests packets. This field is enabled when the <i>Source</i> option is set to <i>IP Address</i> .
<i>Interface</i>	The interface to use when sending the Echo requests packets. This field is enabled when the <i>Source</i> option is set to <i>Interface</i> .
<i>Results</i>	The results of the ping test, which includes information about the reply (if any) received from the host.

Click **Submit** to send the specified number of pings. The results are displayed in the *Results* box.

TraceRoute

Use this page to determine the Layer-3 path a packet takes from the device to a specific IP address or hostname. When you initiate the traceroute command by clicking the **Start** button, the device sends a series of traceroute probes toward the destination. The results list the IP address of each Layer-3 device a probe passes through until it reaches its destination – or fails to reach its destination and is discarded. The information you enter on this page is not saved as part of the device configuration.

To access the *TraceRoute* page, click **System > Utilities > TraceRoute** in the navigation menu.

TraceRoute

Traceroute Fields

Field	Description
<i>Host Name or IP Address</i>	The DNS-resolvable hostname or IP address of the system to attempt to reach.
<i>Probes Per Hop</i>	Traceroute works by sending UDP packets with increasing Time-To-Live (TTL) values. Specify the number of probes sent with each TTL.
<i>MaxTTL</i>	The maximum Time-To-Live (TTL). The traceroute terminates after sending probes that can be Layer-3 forwarded this number of times. If the destination is further away, the traceroute will not reach it.
<i>InitTTL</i>	The initial Time-To-Live (TTL). This value controls the maximum number of Layer-3 hops that the first set of probes may travel.

Traceroute Fields (Continued)

Field	Description
<i>MaxFail</i>	The number of consecutive failures that terminate the traceroute. If the device fails to receive a response for this number of consecutive probes, the traceroute terminates.
<i>Interval</i>	The number of seconds to wait between sending probes.
<i>Port</i>	The UDP destination port number to be used in probe packets. The port number should be a port that the target host is not listening on, so that when the probe reaches the destination, it responds with an <i>ICMP Port Unreachable</i> message.
<i>Size</i>	The size of probe payload in bytes.
<i>Source</i>	Select <i>None</i> , <i>IP Address</i> , or <i>Interface</i> as a source.
<i>IP Address</i>	When the selected <i>Source</i> is <i>IP Address</i> , specify the IP address to use as the source interface.
<i>Interface</i>	When the selected <i>Source</i> is <i>Interface</i> , select the physical port to use as the source interface.
<i>Status</i>	The current status of the traceroute, which can be: • Not Started The traceroute has not been initiated since viewing the page. • In Progress The traceroute has been initiated and is running. • Stopped The traceroute was interrupted by clicking Stop. • Done The traceroute has completed, and information about the traceroute is displayed in the <i>Results</i> area.
<i>Results</i>	Displays the results of the traceroute.

Use the buttons to perform the following tasks:

- Click **Start** to initiate the traceroute.
- Click **Stop** to interrupt the running traceroute.
- Click **Refresh** to refresh the data on the screen with the present state of the data in the switch.

IP Address Conflict Detection

Use the *IP Address Conflict Detection* page to determine whether the IP address configured on the device is the same as the IP address of another device on the same LAN (or on the Internet, for a routable IP address) and to help you resolve any existing conflicts. An IP address conflict can make both this system and the system with the same IP address unusable for network operation.

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



IP Address Conflict Detection

IP Address Conflict Detection Fields

Field	Description
<i>IP Address Conflict Currently Exists</i>	Indicates whether a conflicting IP address has been detected since this status was last reset. False No conflict detected (the subsequent fields on this page display as N/A). True Conflict was detected (the subsequent fields on this page show the relevant information).
<i>Last Conflicting IP Address</i>	The device interface IP address that is in conflict. If multiple conflicts were detected, only the most recent occurrence is displayed.
<i>Last Conflicting MAC Address</i>	The MAC address of the remote host associated with the IP address that is in conflict. If multiple conflicts are detected, only the most recent occurrence is displayed.
<i>Time Since Conflict Detected</i>	The elapsed time (displayed in days, hours, minutes, and seconds) since the last address conflict was detected, provided that you have not yet clicked Clear History .

Use the buttons to perform the following tasks:

- Click **Run Detection** to activate the IP address conflict detection operation in the system.
- Click **Clear History** to reset the IP address conflict detection status information that was last seen by the device.
- Click **Refresh** to refresh the data on the screen with the present state of the data in the switch.

File Transfer

Use the *File Transfer* page to upload files from the device to a remote system and to download files from a remote system to the device.

To access the *File Transfer* page, click **System** > **Utilities** > **Transfer** in the navigation menu.



File Transfer

File Transfer Fields

Field	Description
<i>Transfer Protocol</i>	The protocol to use to transfer the file. Files can be transferred from the device to a remote system using <i>TFTP</i> or <i>FTP</i> . Files can be transferred from a remote system to the device using <i>HTTP</i> , <i>TFTP</i> , or <i>FTP</i> .
<i>Upload</i>	To transfer a file from the device to a remote system using TFTP or FTP, click  in the same row as the desired transfer protocol. The <i>File Upload</i> window appears. Configure the information for the file transfer (described below), and click Begin Transfer to begin the transfer.
<i>Download</i>	To transfer a file from a remote system to the device using HTTP, TFTP or FTP, click  in the same row as the desired transfer protocol. The <i>File Download</i> window appears. Configure the information for the file transfer (described below), and click Begin Transfer to begin the transfer.

Uploading Files

When you click , the *File Upload* window appears. The following information describes the fields in the *File Upload* window for all protocols.

File Upload Fields

Field	Description
<i>File Type</i>	Specify the type of file to transfer from the device to a remote system. Code Select this option to transfer an image. Configuration Select this option to transfer a copy of the stored configuration file (startup-config) to a remote system. Backup Configuration Select this option to transfer a copy of the stored backup configuration (backup-config) from the device to a remote system. Script File Select this option to transfer a custom text configuration script from the device to a remote system. CLI Banner Select this option to transfer the file containing the text to be displayed on the CLI before the login prompt to a remote system. Crash Log Select this option to transfer the system crash log to a remote system. Operational Log Select this option to transfer the system operational log to a remote system. Startup Log Select this option to transfer the system startup log to a remote system. Trap Log Select this option to transfer the system trap records to a remote system. Factory Defaults Select this option to transfer the factory default configuration file to a remote system. Error Log Select this option to transfer the system error (persistent) log, which is also known as the event log, to a remote system. Buffered Log Select this option to transfer the system buffered (in-memory) log to a remote system.
<i>Image</i>	If the selected <i>File Type</i> is <i>Code</i> , specify whether to transfer the <i>Active</i> or <i>Backup</i> image to a remote system.
<i>Server Address</i>	Specify the IPv4 address, IPv6 address, or DNS-resolvable hostname of the remote server that will receive the file.
<i>File Path</i>	Specify the path on the server where you want to put the file.
<i>File Name</i>	Specify the name that the file will have on the remote server.
<i>User Name</i>	For FTP transfers, if the server requires authentication, specify the user name for remote login to the server that will receive the file.
<i>Password</i>	For FTP transfers, if the server requires authentication, specify the password for remote login to the server that will receive the file.
<i>Progress</i>	Represents the completion percentage of the file transfer. The file transfer begins after you complete the required fields and click  to the right of this field.
<i>Digital Signature Verification</i>	For <i>Code</i> and <i>Configuration</i> file types this option, when checked, will verify the file download with the digital signature.
<i>Status</i>	Provides information about the status of the file transfer.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Downloading Files

When you click , the *File Download* window appears. The following information describes the fields in the *File Download* window for all protocols.

File Download Fields

Field	Description
<i>File Type</i>	Specify the type of file to transfer to the device: • Code Select this option to transfer a new image to the device. The code file is stored as the backup image. • Configuration Select this option to update the stored configuration file (startup-config). If the file has errors, the update will be stopped. • Script File Select this option to transfer a text-based configuration script to the device. You must use the command-line interface (CLI) to validate and activate the script. • CLI Banner Select this option to transfer the CLI banner file to the device. This file contains the text to be displayed on the CLI before the login prompt. • IAS Users Select this option to transfer an Internal Authentication Server (IAS) users database file to the device. The IAS user database stores a list of user name and (optional) password values for local port-based user authentication. • SSH-1 RSA Key File Select this option to transfer an SSH-1 Rivest-Shamir-Adleman (RSA) key file to the device. SSH key files contain information to authenticate SSH sessions for remote CLI-based access to the device. • SSH-2 RSA Key PEM File Select this option to transfer an SSH-2 Rivest-Shamir-Adleman (RSA) key file (PEM Encoded) to the device. • SSH-2 DSA Key PEM File Select this option to transfer an SSH-2 Digital Signature Algorithm (DSA) key file (PEM Encoded) to the device. • SSL Trusted Root Certificate PEM File Select this option to transfer an SSL Trusted Root Certificate file (PEM Encoded) to the device. SSL files contain information to encrypt, authenticate, and validate HTTPS sessions. • SSL Server Certificate PEM File Select this option to transfer an SSL Server Certificate file (PEM Encoded) to the device. • SSL DH Weak Encryption Parameter PEM File Select this option to transfer an SSL Diffie-Hellman Weak Encryption Parameter file (PEM Encoded) to the device. • SSL DH Strong Encryption Parameter PEM File Select this option to transfer an SSL Diffie-Hellman Strong Encryption Parameter file (PEM Encoded) to the device. Note: • To download SSH key files, SSH must be administratively disabled, and there can be no active SSH sessions. • To download SSL related files, HTTPS must be administratively disabled.
<i>Select File</i>	If the <i>Transfer Protocol</i> is set to <i>HTTP</i> , browse to the directory where the file is located and select the file to transfer to the device. This field is not present if the <i>Transfer Protocol</i> is TFTP or FTP .
<i>Server Address</i>	For TFTP or FTP transfers, specify the IPv4 address, IPv6 address, or DNS-resolvable hostname of the remote server.
<i>File Path</i>	For TFTP or FTP transfers, specify the path on the server where the file is located.
<i>File Name</i>	For TFTP or FTP transfers, specify the name of the file you want to transfer to the device.
<i>User Name</i>	For FTP transfers, if the server requires authentication, specify the user name for remote login to the server where the file resides.
<i>Password</i>	For FTP transfers, if the server requires authentication, specify the password for remote login to the server that will receive the file.
<i>Progress</i>	Represents the completion percentage of the file transfer. The file transfer begins after you complete the required fields and click  to the right of this field.
<i>Digital Signature Verification</i>	For <i>Code</i> and <i>Configuration</i> file types this option, when checked, will verify the file download with the digital signature.
<i>Status</i>	Provides information about the status of the file transfer.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

AutoInstall

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

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

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

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

To display the *AutoInstall Configuration* page, click **System > Firmware > AutoInstall**.

AutoInstall Configuration	
Admin Mode	☒ Start ☐ Stop
Persistent Mode	☐
AutoSave Mode	☐
AutoReboot Mode	☒
Retry Count	3 (1 to 31)
Status	AutoInstall is completed.
Submit Refresh Cancel	

© Copyright 2013-2014 Ubiquiti Networks, Inc.

AutoInstall Configuration

Autoinstall Configuration Fields

Field	Description
<i>Admin Mode</i>	The current administrative mode of the AutoInstall feature: • Start AutoInstall is enabled, and the feature will attempt to automatically configure the device during the next boot cycle. • Stop AutoInstall is disabled. The automatic process will begin only if no configuration file is located during the next boot cycle.
<i>Persistent Mode</i>	If this option is selected, the settings you configure on this page are automatically saved to persistent memory in the startup-config file when you apply the changes. If this option is cleared, the device treats these settings like any other applied changes (i.e., the changes are not retained across a reboot unless you save the configuration).
<i>AutoSave Mode</i>	If this option is selected, the downloaded configuration is automatically saved to persistent storage. If this option is cleared, you must explicitly save the downloaded configuration in non-volatile memory for the configuration to be available for the next reboot.
<i>AutoReboot Mode</i>	If this option is selected, the switch automatically reboots after a new image is successfully downloaded and makes the downloaded image the active image. If this option is cleared, the device continues to boot with the current image. The downloaded image will not become the active image until the device reboots.
<i>Retry Count</i>	When attempting to retrieve the DHCP-specified configuration file, this value represents the number of times the TFTP client on the device tries to use unicast requests before reverting to broadcast requests.
<i>Status</i>	The current status of the AutoInstall process.

Use the buttons to perform the following tasks:

- If you change any settings on this page, click **Submit** to apply the changes.
- To reset the fields to their original values, click **Cancel**.
- Click **Refresh** to display the most recently configured AutoInstall state from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.


Managing SNMP Traps

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

System Trap Log

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

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



System Trap Log

System Trap Log Fields

Field	Description
<i>Trap Log Capacity</i>	The maximum number of traps stored in the log. If the number of traps exceeds the capacity, the entries will overwrite the oldest entries.
<i>Number of Traps Since Last Reset</i>	The number of traps generated since the trap log entries were last cleared.
<i>Number of Traps Since Log Last Viewed</i>	The number of traps that have occurred since the traps were last displayed. Displaying the traps by any method (terminal interface display, web display, upload file from switch, etc.) will cause this counter to be cleared to 0.
<i>Log</i>	The sequence number of this trap.
<i>System Up Time</i>	The time at which this trap occurred, expressed in days, hours, minutes and seconds since the last reboot of the switch.
<i>Trap</i>	Displays the information identifying the trap.


Use the buttons to perform the following tasks:

- Click **Clear Log** to clear all entries in the log. Subsequent displays of the log will only show new log entries.
- Click **Refresh** to refresh the data on the screen with the present state of the data in the switch.

System Trap Flags

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

To access the page, click **System > Advanced Configuration > Trap Manager > Trap Flags** page.



System Trap Flags

The fields available on the *System Trap Flags* page depends on the packages installed on your system. For example, if your system does not have the BGP4 package installed, the *BGP Traps* field is not available. The illustration above and the table below show the fields that are available on a system with all packages installed.

System Trap Flags Fields

Field	Description
<i>Authentication</i>	When selected, this option enables activation of authentication failure traps by selecting the corresponding line on the pulldown entry field. This feature is enabled by default.
<i>Link Up/Down</i>	When selected, this option enables activation of link status traps by selecting the corresponding line on the pulldown entry field. This feature is enabled by default.
<i>Multiple Users</i>	When selected, this option enables activation of multiple user traps by selecting the corresponding line on the pulldown entry field. This feature is enabled by default. This trap is triggered when the same user ID is logged into the switch more than once at the same time (either via Telnet or the serial port).
<i>Spanning Tree</i>	When selected, this option enables activation of spanning tree traps by selecting the corresponding line on the pulldown entry field. This feature is enabled by default.
<i>ACL Traps</i>	When selected, this option enables activation of ACL traps by selecting the corresponding line on the pulldown entry field. This feature is disabled by default.
<i>Power Supply Module State</i>	When selected, this option enables SNMP notifications when power supply events occur.
<i>Temperature</i>	When selected, this option enables SNMP notifications when temperature events occur.

Use the buttons to perform the following tasks:

- If you make any changes to this page, click **Submit** to apply the changes to the system.
- Click **Refresh** to refresh the data on the screen with the present state of the data in the switch.

Managing the DHCP Server

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

DHCP Server Global Configuration

Use the *DHCP Server Global Configuration* page to configure DHCP global parameters.

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

DHCP Server Global Configuration

DHCP Server Global Configuration Fields

Field	Description
<i>Admin Mode</i>	Used to <i>Enable</i> or <i>Disable</i> the DHCP server administrative mode. When enabled, the device can be configured to automatically allocate TCP/IP configurations for clients.
<i>Conflict Logging Mode</i>	Used to <i>Enable</i> or <i>Disable</i> the logging mode for IP address conflicts. When enabled, the system stores information IP address conflicts that are detected by the DHCP server.
<i>Bootp Automatic Mode</i>	Used to <i>Enable</i> or <i>Disable</i> the BOOTP automatic mode. When enabled, the DHCP server supports the allocation of automatic addresses for BOOTP clients. When disabled the DHCP server supports only static addresses for BOOTP clients.
<i>Ping Packet Count</i>	The number of packets the server sends to a pool address to check for duplication as part of a ping operation. If the server receives a response to the ping, the address is considered to be in conflict and is removed from the pool.

Use the buttons to perform the following tasks:

- If you change any settings on this page, click **Submit** to apply the changes.
- To reset the fields to their original values, click **Cancel**.
- Click **Refresh** to refresh the data on the screen with the present state of the data in the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

DHCP Server Pool Configuration

Use the *DHCP Server Pool Configuration* page to create the pools of addresses that can be assigned by the server.

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

DHCP Server Pool Configuration

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

DHCP Server Pool Configuration Fields

Field	Description
<i>Pool Name</i>	Select the pool to configure. The menu includes all pools that have been configured on the device.
<i>Type of Binding</i>	Specifies the type of binding for the pool. The options are: Manual You statically assign an IP address to a client based on the client's MAC address. Dynamic The DHCP server can assign the client any available IP address within the pool. This type is also known as Automatic.
<i>Network Base Address</i>	Dynamic pools only – The network portion of the IP address. A DHCP client can be offered any available IP address within the defined network as long as it has not been configured as an excluded address.
<i>Network Mask</i>	Dynamic pools only – The subnet mask associated with the Network Base Address that separates the network bits from the host bits.

DHCP Server Pool Configuration Fields (Continued)

Field	Description
<i>Client Name</i>	Manual pools only – The system name of the client. The Client Name should not include the domain name. This field is optional.
<i>Hardware Address Type</i>	Manual pools only – The protocol type (<i>Ethernet</i> [default] or <i>IEEE802</i>) used by the client's hardware platform. This value is used in response to requests from BOOTP clients.
<i>Hardware Address</i>	Manual pools only – The MAC address of the DHCP client.
<i>Client ID</i>	Manual pools only – The value some DHCP clients send in the <i>Client Identifier</i> field of DHCP messages. This value is typically identical to the <i>Hardware Address</i> value. In some systems, such as Microsoft DHCP clients, the client identifier is required instead of the hardware address. If the client's DHCP request includes the client identifier, the <i>Client ID</i> field on the DHCP server must contain the same value, and the <i>Hardware Address Type</i> field must be set to the appropriate value. Otherwise, the DHCP server will not respond to the client's request.
<i>Host IP Address</i>	Manual pools only – The IP address to offer the client.
<i>Host Mask</i>	Manual pools only – This field specifies the subnet mask to be statically assigned to a DHCP client.
<i>Lease Expiration</i>	Indicates whether the information the server provides to the client should expire. Enable Allows the lease to expire. If you select this option, you can specify the amount of time the lease is valid in the <i>Lease Duration</i> field. Disable Sets an infinite lease time. For <i>Dynamic</i> bindings, an infinite lease time implies a lease period of 60 days. For a <i>Manual</i> binding, an infinite lease period never expires.
<i>Lease Duration</i>	The number of <i>Days</i> , <i>Hours</i> , and <i>Minutes</i> the lease is valid. This field cannot be configured if the <i>Lease Expiration</i> is disabled.
<i>Next Server Address</i>	The IP address of the next server the client should contact in the boot process. For example, the client might be required to contact a TFTP server to download a new image file. Use the buttons as follows:  Click this button to configure the <i>Next Server Address</i> field.  Click this button to reset the field to the default value.
<i>Default Router, DNS Server, NetBIOS Server</i> – To configure settings for one or more default routers, DNS servers, or NetBIOS servers that can be used by the client(s) in the pool, use the buttons available in the appropriate table to perform the following tasks:  To add an entry to the server list, click this button and enter the IP address of the server to add.  To edit the address of a configured server, click this button associated with the entry to edit and update the address.  To delete an entry from the list, click this button associated with the entry to remove. To delete <i>all</i> entries from the list, click this button in the heading row.	
<i>Default Router</i>	Lists the IP address of each router to which the client(s) in the pool should send traffic. The default router should be in the same subnet as the client.
<i>DNS Server</i>	Lists the IP address of each DNS server the client(s) in the pool can contact to perform address resolution.
<i>NetBIOS Server</i>	Lists the IP address of each NetBIOS Windows Internet Naming Service (WINS) name server that is available for the selected pool.

Use the buttons to perform the following tasks:

- After you configure values for the DHCP address pool, click **Submit** to create the pool and apply the changes to the system.
- To reset the fields to their original values, click **Cancel**.
- To delete a pool, select the pool from the *Pool Name* drop-down menu and click **Delete**.
- Click **Refresh** to refresh the data on the screen with the present state of the data in the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

DHCP Server Pool Options

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

To access the page, click **System > Advanced Configuration > DHCP Server > Pool Options** in the navigation menu. The page displays the fields shown below only if DHCP pools are configured on the system.

DHCP Server Pool Options

DHCP Server Pool Options Fields

Field	Description
<i>Pool Name</i>	Select the DHCP pool to view or configure. The menu lists all pools that are configured on the switch.
<i>NetBIOS Node Type</i>	The method the client should use to resolve NetBIOS names to IP addresses. The options are: B-Node Broadcast Broadcast only P-Node Peer-to-Peer NetBIOS name server only M-Node Mixed Broadcast, then NetBIOS name server H-Node Hybrid NetBIOS name server, then broadcast Use the buttons as follows: - Click this button to configure the field. - Click this button to reset the field to the default value.
<i>Domain Name</i>	The default domain name to configure for all clients in the selected pool. Use the buttons as follows: - Click this button to configure the field. - Click this button to reset the field to the default value.
<i>Bootfile Name</i>	The name of the default boot image that the client should attempt to download from a specified boot server. Use the buttons as follows: - Click this button to configure the field. - Click this button to reset the field to the default value.
The lower section of the page contains the option table which shows the Vendor Options that have been added to the selected pool.	
<i>Option Name</i>	Identifies whether the entry is a fixed option or a vendor-defined option (<i>Vendor</i>).
<i>Option Code</i>	The number that uniquely identifies the option.
<i>Option Type</i>	Specifies the type of option associated with the option code configured for the selected pool: ASCII The option type is a text string. HEX The option type is a hexadecimal number. IP Address The option type is an IP address.
<i>Option Value</i>	The data associated with the <i>Option Code</i>. When adding or editing a vendor option, the field(s) available for configuring the value depend on the selected <i>Option Type</i>. If the value you configure contains invalid characters for the selected <i>Option Type</i>, the configuration cannot be applied.

Use the buttons to perform the following tasks:

- To add a vendor option, click **Add Vendor Option**, configure the available fields, and click **Submit** to apply the changes.
- To edit a vendor option, select the entry to change and click **Edit**. Change the settings as needed (*Pool Name* and *Option Code* are not configurable) and click **Submit** to apply the changes.
- To remove a vendor option, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.
- Click **Refresh** to refresh the data on the screen with the present state of the data in the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

DHCP Server Bindings Information

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

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



DHCP Server Bindings

DHCP Server Bindings Fields

Field	Description
<i>IP Address</i>	The IP Address of the DHCP client.
<i>Hardware Address</i>	The MAC address of the DHCP client.
<i>Lease Time Left</i>	The amount of time left until the lease expires in days, hours, and minutes.
<i>Pool Allocation Type</i>	The type of binding used: • Dynamic The address was allocated dynamically from a pool that includes a range of IP addresses. • Manual A static IP address was assigned based on the MAC address of the client. • Inactive The pool is not in use.

If you change any settings, click **Submit** to apply the changes to the system.

- To remove an entry from the table, select each entry to delete and click **Clear Entries**. You must confirm the action before the binding is deleted.
- Click **Refresh** to refresh the data on the screen with the present state of the data in the switch.

DHCP Server Statistics

Use the *DHCP Server Statistics* page to view information about the DHCP server bindings and messages. To access the page, click **System > Advanced Configuration > DHCP Server > Statistics** in the navigation menu.

DHCP Server Statistics	
Automatic Bindings	0
Expired Bindings	0
Malformed Messages	0
Messages Received	
DHCPDISCOVER	0
DHCPREQUEST	0
DHCPDECLINE	0
DHCPRELEASE	0
dhcpinform	0
Messages Sent	
DHCPOFFER	0
DHCPACK	0
DHCPNAK	0

DHCP Server Statistics

DHCP Server Statistics Fields

Field	Description
<i>Automatic Bindings</i>	Shows the number of automatic bindings on the DHCP server.
<i>Expired Bindings</i>	Shows the number of expired bindings on the DHCP server.
<i>Malformed Messages</i>	Shows the number of the malformed messages.
<i>Message Received section:</i>	
<i>DHCPDISCOVER</i>	Shows the number of DHCPDISCOVER messages received by the DHCP server.
<i>DHCPREQUEST</i>	Shows the number of DHCPREQUEST messages received by the DHCP server.
<i>DHCPDECLINE</i>	Shows the number of DHCPDECLINE messages received by the DHCP server.
<i>DHCPRELEASE</i>	Shows the number of DHCPRELEASE messages received by the DHCP server.
<i>DHCPINFORM</i>	Shows the number of DHCPINFORM messages received by the DHCP server.
<i>DHCPOFFER</i>	Shows the number of DHCPOFFER messages sent by the DHCP server.
<i>DHCPACK</i>	Shows the number of DHCPACK messages sent by the DHCP server.
<i>DHCPNAK</i>	Shows the number of DHCPNAK messages sent by the DHCP server.
<i>Message Sent section:</i>	
<i>DHCPOFFER</i>	The number of DHCP offer messages the DHCP server has sent to DHCP clients in response to DHCP discovery messages it has received.
<i>DHCPACK</i>	The number of DHCP acknowledgement messages the DHCP server has sent to DHCP clients in response to DHCP request messages. The server sends this message after a client accepts the server's offer. The message includes information about the lease time and any other configuration information that the DHCP client has requested.
<i>DHCPNAK</i>	The number of negative DHCP acknowledgement messages the DHCP server has sent to DHCP clients. This type of message is sent if the client requests an IP address already in use or if the server does not renew the lease.

Use the buttons to perform the following tasks:

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

DHCP Server Conflicts Information

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

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



DHCP Server Conflicts Information

DHCP Server Conflicts Information Fields

Field	Description
<i>IP Address</i>	The IP address that has been detected as a duplicate.
<i>Detection Method</i>	The method used to detect the conflict, which is one of the following: • Gratuitous ARP The DHCP client detected the conflict by broadcasting an ARP request to the address specified in the DHCP offer message sent by the server. If the client receives a reply to the ARP request, it declines the offer and reports the conflict. • Ping The server detected the conflict by sending an ICMP echo message (ping) to the IP address before offering it to the DHCP client. If the server receives a response to the ping, the address is considered to be in conflict and is removed from the pool. • Host Declined The server received a DHCPDECLINE message from the host. A DHCPDECLINE message indicates that the host has discovered that the IP address is already in use on the network.
<i>Detection Time</i>	The time when the conflict was detected in days, hours, minutes, and seconds since the system was last reset (i.e., system up time).

Use the buttons to perform the following tasks:

- Click **Clear Entries** to clear all of the address conflict entries.
- Click **Refresh** to update the information on the screen.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Configuring Time Ranges

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

Time Range Configuration

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

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

Time Range Summary

Time Range Summary Fields

Field	Description
<i>Admin Mode</i>	Used to <i>Enable</i> or <i>Disable</i> the Time Range administrative mode. When enabled, actions with subscribed components are performed for existing time range entries.
<i>Time Range Name</i>	The unique ID or name that identifies this time range. A time-based ACL rule can reference the name configured in this field.
<i>Time Range Status</i>	Shows whether the time range is <i>Active</i> or <i>Inactive</i> . A time range is <i>Inactive</i> if the current day and time do not fall within any time range entries configured for the time range.
<i>Periodic Entry Count</i>	The number of periodic time range entries currently configured for the time range.
<i>Absolute Entry</i>	Shows whether an absolute time entry is currently configured for the time range.

Use the buttons to perform the following tasks:

- To add a time range, click **Add**, enter a name for the time range configuration, and click **Submit** to create the time range.
- To delete a configured time range, select each entry to delete, click **Remove**, and confirm the action.
- If you change the *Admin Mode* setting on this page, click **Submit** to apply the change.
- Click **Refresh** to update the information on the screen.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Time Range Entry Configuration

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



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

To access this page, click **System > Advanced Configuration > Time Ranges > Entry Configuration**.

Time Range Entry Summary

Time Range Entry Summary Fields

Field	Description
<i>Time Range Name</i>	Lists the available time ranges or blank if no time ranges have been defined yet.
<i>Entry Type</i>	The type of time range entry, which is one of the following: Absolute Occurs once or has an undefined start or end period. The duration of an absolute entry can be hours, days, or even years. Each time entry configuration can have only one absolute entry. Periodic Recurring entry that takes place at fixed intervals. This type of entry occurs at the same time on one or more days of the week.
<i>Starts</i>	For an absolute entry, indicates the time, day, month, and year that the entry begins. If this field is blank, the absolute entry became active when it was configured. For a periodic entry, indicates the time and day(s) of the week that the entry begins.
<i>Ends</i>	For an absolute entry, indicates the time, day, month, and year that the entry ends. If this field is blank, the absolute entry does not have a defined end. For a periodic entry, indicates the time and day(s) of the week that the entry ends.
<i>Add Absolute Time Range</i> dialog box – When you click Add Absolute , this dialog box appears with the following fields:	
<i>Time Range Name</i>	The time range configuration that will include the absolute time range entry.
<i>Start Time</i>	Select this option to configure values for the <i>Start Date</i> and the <i>Starting Time of Day</i> . If this option is not selected, the entry becomes active immediately.
<i>Start Date</i>	Click to select the day, month, and year when this entry becomes active. This field can be configured only if the <i>Start Time</i> option is selected.
<i>Starting Time of Day</i>	Specify the time of day that the entry becomes active by entering the information in the field or by using the scroll bar in the <i>Choose Time</i> pop-up window. Click Now to use the current time of day. Click Done to close the <i>Choose Time</i> window. This field can be configured only if the <i>Start Time</i> option is selected.

Time Range Entry Summary Fields (Continued)

Field	Description
<i>End Time</i>	Select this option to configure values for the <i>End Date</i> and the <i>Ending Time of Day</i> . If this option is not selected, the entry does not have an end time; after the configured <i>Start Time</i> begins, the entry will remain active indefinitely.
<i>End Date</i>	Click  to select the day, month, and year when this entry should no longer be active. This field can be configured only if the <i>End Time</i> option is selected.
<i>Ending Time of Day</i>	Specify the time of day that the entry becomes inactive by entering the information in the field or by using the scroll bar in the <i>Choose Time</i> pop-up window. Click Now to use the current time of day. Click Done to close the <i>Choose Time</i> pop-up window. This field can be configured only if the <i>End Time</i> option is selected.
Add Periodic Time Range dialog box – When you click Add Periodic , this dialog box appears, with the following fields:	
<i>Time Range Name</i>	The time range configuration that will include the Periodic time range entry.
<i>Applicable Days</i>	Select the days on which the Periodic time range entry is active: • Daily Every day of the week • Weekdays Monday through Friday • Weekend Saturday and Sunday • Days of Week User-defined start days
<i>Start Days</i>	Indicates on which days the time entry becomes active. If the selected option in the <i>Applicable Days</i> field is <i>Days of Week</i> , select one or more days on which the entry becomes active. To select multiple days, press and hold CTRL and select each desired start day.
<i>Starting Time of Day</i>	Specify the time of day that the entry becomes active by entering the information in the field or by using the scroll bar in the <i>Choose Time</i> pop-up window. Click Now to use the current time of day. Click Done to close the <i>Choose Time</i> pop-up window.
<i>End Days</i>	Indicates on which days the time entry ends. If the selected option in the <i>Applicable Days</i> field is <i>Days of Week</i> , select one or more days on which the entry ends. To select multiple days, press and hold CTRL and select each desired end day.
<i>Ending Time of Day</i>	Specify the time of day that the entry becomes inactive by entering the information in the field or by using the scroll bar in the <i>Choose Time</i> pop-up window. Click Now to use the current time of day. Click Done to close the <i>Choose Time</i> pop-up window.

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

- To add an absolute time range entry, click **Add Absolute**, configure the settings to define the absolute time range, and then click **Submit** to apply the changes. If the **Add Absolute** button is not available, an absolute entry already exists for the time range specified by *Time Range Name*.
- To add a periodic time range entry, click **Add Periodic** and specify the days and times that the entry is in effect.
- To delete a time range entry, select each entry to delete, click **Remove**, and confirm the action.
- Click **Refresh** to update the information on the screen.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Configuring DNS

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

DNS Global Configuration

Use the *DNS Global Configuration* page to configure global DNS settings and to view DNS client status information. To access this page, click **System > Advanced Configuration > DNS > Configuration**.

DNS Global Configuration

DNS Global Configuration Fields

Field	Description
<i>Admin Mode</i>	The administrative mode, <i>Enable</i> or <i>Disable</i> (default), of the DNS client.
<i>Default Domain Name</i>	The default domain name (255 characters maximum) that the DNS client uses to complete unqualified host names. After a default domain name is configured (default: not configured), a host name entered without domain name information is appended with the default domain name. For example, if the default domain name is <i>.com</i> and the user enters hotmail as the host name, then the host name is changed to <i>hotmail.com</i> .
<i>Retry Number</i>	The number of times to resend DNS queries to a DNS server on the network. Range is 0 to 100. The default is 2.
<i>Response Timeout</i>	The number of seconds to allow a DNS server to respond to a request before a retry. Range is 0 to 3600. Default is 3.
<i>Domain List</i>	The domain names that have added to the DNS client's domain list. If a DNS query that includes the default domain name is not resolved, the DNS client uses these domain names, in the order they appear in this list, to extend the hostname into a fully-qualified domain name. Use the buttons as follows: - To create a new list of domain names, click this button, enter the name of the list, and click Submit. Repeat this step to add multiple domains to the default domain list. - To remove a domain from the domain list, click this button and then confirm the action.
<i>DNS Server</i>	A unique IPv4 or IPv6 address used to identify a DNS server. The order in which you add servers determines the precedence of the server. The DNS server that you add first has the highest precedence and will be used before other DNS servers that you add. Use the buttons as follows: - Click this button to configure the associated DNS server. - To delete the associated DNS server entry, Click this button and then confirm the action.


Use the buttons to perform the following tasks:

- If you change any settings on this page, click **Submit** to apply the changes.
- Click **Refresh** to update the information on the screen.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

DNS IP Mapping Configuration

Use the *DNS IP Mapping* page to view and manage the Static and Dynamic entries in the DNS IP mapping table. To access this page, click **System > Advanced Configuration > DNS > IP Mapping** in the menu.



DNS IP Mapping

DNS IP Mapping Fields

Field	Description
<i>Entry Type</i>	Type of DNS entry: Static An entry that has been manually configured on the device. Dynamic An entry that the device has learned by using a configured DNS server to resolve a hostname.
<i>Host Name</i>	The name that identifies the system. For <i>Static</i> entries, specify the <i>Host Name</i> after you click Add . A host name can contain up to 255 characters if it contains multiple levels in the domain hierarchy, but each level (the portion preceding a period) can contain a maximum of 63 characters. If the host name you specify is a single level (does not contain any periods), the maximum number of allowed characters is 63.
<i>IP Address</i>	The IPv4 or IPv6 address associated with the configured <i>Host Name</i> . For <i>Static</i> entries, specify the <i>IP Address</i> after you click Add . You can specify either an IPv4 or an IPv6 address.
<i>Dynamic Entry fields</i> – The following fields include values for Dynamic entries only. For Static entries, these fields are blank.	
<i>Total Time</i>	The number of seconds that the entry will remain in the table.
<i>Elapsed Time</i>	The number of seconds that have passed since the entry was added to the table. When the <i>Elapsed Time</i> reaches the <i>Total Time</i> , the entry times out and is removed from the table.
<i>Dynamic Type</i>	The type of address in the entry; for example IP, or X.121 (less common).

Use the buttons to perform the following tasks:

- To statically map an IP address to a hostname, click **Add**, configure the *Host Name* and *IP Address* fields in the *Add DNS Entry* dialog box, and then click **Submit** to apply the changes.
- To delete one or more entries, select each entry to delete, click **Remove**, and confirm the deletion.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

DNS Source Interface Configuration

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

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

DNS Source Interface Configuration

DNS Source Interface Configuration Fields

Field	Description
<i>Type</i>	The type of interface to use as the source interface: None The primary IP address of the originating (outbound) interface is used as the source address. Interface The primary IP address of a physical port is used as the source address. VLAN The primary IP address of a VLAN routing interface is used as the source address. Tunnel The primary IP address of a tunnel interface is used as the source address.
<i>Interface</i>	When the selected <i>Type</i> is <i>Interface</i> , select the physical port to use as the source interface.
<i>VLAN</i>	When the selected <i>Type</i> is <i>VLAN</i> , select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.
<i>Tunnel</i>	When the selected <i>Type</i> is <i>Tunnel</i> , select the tunnel interface to use as the source interface.

Use the buttons to perform the following tasks:

- If you change any of the settings on the page, click **Submit** to apply the changes to system.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Configuring SNTP Settings

The EdgeSwitch software supports the Simple Network Time Protocol (SNTP). SNTP assures accurate network device clock time synchronization up to the millisecond. Time synchronization is performed by a network SNTP server. The EdgeSwitch software operates only as an SNTP client and cannot provide time services to other systems.

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

The following is an example of stratum:

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

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

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

T1: Time at which the original request was sent by the client.

T2: Time at which the original request was received by the server.

T3: Time at which the server sent a reply.

T4: Time at which the client received the server's reply.

The device can poll Unicast and Broadcast server types for the server time.

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

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

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

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

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

SNTP Global Configuration

Use the *SNTP Global Configuration* page to view and adjust SNTP parameters. To display the page, click **System > Advanced Configuration > SNTP > Global Configuration** in the navigation menu.

The screenshot shows the 'SNTP Global Configuration' page. At the top, there are five tabs: 'Global Configuration' (selected), 'Global Status', 'Server Configuration', 'Server Status', and 'Source Interface Configuration'. Below the tabs, the title 'SNTP Global Configuration' is displayed. The main area contains several configuration fields: 'Client Mode' is a dropdown menu set to 'Unicast'; 'Port' is a text box with 'None' and two buttons (a gear icon and a reset icon); 'Unicast Poll Interval (Seconds)' is a text box with '6' and a range of '05 to 100'; 'Broadcast Poll Interval (Seconds)' is a text box with '6' and a range of '05 to 100'; 'Unicast Poll Timeout (Seconds)' is a text box with '5' and a range of '(1 to 30)'; 'Unicast Poll Retry' is a text box with '1' and a range of '(0 to 10)'; and 'Number of Servers Configured' is a text box with '2'. At the bottom, there are three buttons: 'Submit', 'Refresh', and 'Cancel'. A copyright notice '© Copyright 2013-2014 Ubiquiti Networks, Inc.' is visible in the bottom right corner.

SNTP Global Configuration

SNTP Global Configuration Fields

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


Use the buttons to perform the following tasks:

- If you change any of the settings on the page, click **Submit** to apply the changes to system.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

SNTP Global Status

Use the *SNTP Global Status* page to view information about the system's SNTP client. To access the page, click **System > Advanced Configuration > SNTP > Global Status** in the navigation menu.



SNTP Global Status

SNTP Global Status Fields

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

SNTP Global Status Fields (Continued)

Field	Description
<i>Unicast Sever Max Entries</i>	The maximum number of unicast server entries that can be configured on this client.
<i>Unicast Server Current Entries</i>	The number of current valid unicast server entries configured for this client.
<i>Broadcast Count</i>	The number of unsolicited broadcast SNTP messages that have been received and processed by the SNTP client since last reboot.

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

SNTP Server Configuration

Use the *SNTP Server Configuration* page to view and modify information for adding and modifying Simple Network Time Protocol SNTP servers. To display the *SNTP Server Configuration* page, click **System > Advanced Configuration > SNTP > Server Configuration** in the navigation menu.



SNTP Server Configuration

SNTP Server Configuration Fields

Field	Description
<i>SNTP Server</i>	Select the IP address of a user-defined SNTP server to view or modify information about an SNTP server, or click Add to configure a new SNTP server. You can define up to three SNTP servers.
<i>Type</i>	Select IPv4 if you entered an IPv4 address, DNS if you entered a hostname.
<i>Port</i>	Enter a port number from 1 to 65535. The default is 123.
<i>Priority</i>	Enter a priority from 1 to 3, with 1 being the highest priority. The switch will attempt to use the highest priority server and, if it is not available, will use the next highest server.
<i>Version</i>	Enter the protocol version number.
<i>Add SNTP Server Dialog Box</i> – When you click Add , this dialog box appears, containing the following additional field:	
<i>Host Name or IP Address</i>	Specify the IPv4 address, IPv6 address, or DNS-resolvable host name of the SNTP server. Unicast SNTP requests will be sent to this address. The address you enter is displayed in the SNTP Server field on the main page. The address type is automatically detected.

Use the buttons to perform the following tasks:

- To add an SNTP server, click **Add**, configure the fields as needed, and click **Submit** to apply the changes. The SNTP server is added and appears in the *SNTP Server* list.
- To modify settings for an existing SNTP server, select the entry to update, click **Edit**, update the fields as needed, and click **Submit** to apply the changes. You cannot edit the host name or address of a server.
- To remove an SNTP server from the list, select each entry to delete, click **Remove**, and confirm the deletion.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

SNTP Server Status

The *SNTP Server Status* page displays status information about the SNTP servers configured on your switch. To access the *SNTP Server Status* page, click **System > Advanced Configuration > SNTP > Server Status** in the navigation menu.

Address	Last Update Time	Last Attempt Time	Last Attempt Status	Requests	Failed Requests
1.ubnt.pool.ntp.org	Jul 18 23:42:22 201	Jul 19 00:21:03 201	Success	371	1
2.ubnt.pool.ntp.org	Jan 1 00:09:00 197	Jan 1 00:00:00 197	Other	0	0

SNTP Server Status

SNTP Server Status Fields

Field	Description
<i>Address</i>	The existing server addresses. If no server configuration exists, <i>No SNTP server exists</i> is displayed on-screen.
<i>Last Update Time</i>	The local date and time (UTC) that the response from this server was used to update the system clock.
<i>Last Attempt Time</i>	The local date and time (UTC) that this SNTP server was last queried.
<i>Last Attempt Status</i>	The status of the last SNTP request to this server. If no packet has been received from this server, a status of <i>Other</i> is displayed: Other None of the following enumeration values. Success The SNTP operation was successful and the system time was updated. Request Timed Out A directed SNTP request timed out without receiving a response from the SNTP server. Bad Date Encoded The time provided by the SNTP server is not valid. Version Not Supported The SNTP version supported by the server is not compatible with the version supported by the client. Server Unsynchronized The SNTP server is not synchronized with its peers. This is indicated via the 'leap indicator' field on the SNTP message. Server Kiss Of Death The SNTP server indicated that no further queries were to be sent to this server. This is indicated by a stratum field equal to 0 in a message received from a server.
<i>Requests</i>	The number of SNTP requests made to this server since last agent reboot.
<i>Failed Requests</i>	The number of failed SNTP requests made to this server since last reboot.

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

SNTP Source Interface Configuration

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

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

SNTP Source Interface Configuration

SNTP Source Interface Configuration Fields

Field	Description
<i>Type</i>	The type of interface to use as the source interface: None The primary IP address of the originating (outbound) interface is used as the source address. Interface The primary IP address of a physical port is used as the source address. VLAN The primary IP address of a VLAN routing interface is used as the source address. Tunnel The primary IP address of a tunnel interface is used as the source address.
<i>Interface</i>	When the selected <i>Type</i> is <i>Interface</i> , select the physical port to use as the source interface.
<i>VLAN ID</i>	When the selected <i>Type</i> is <i>VLAN</i> , select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.
<i>Tunnel ID</i>	When the selected <i>Type</i> is <i>Tunnel</i> , select the tunnel interface to use as the source interface.

Use the buttons to perform the following tasks:

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

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Configuring the Time Zone

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

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

Time Zone Summary

Time Zone Summary Fields

Field	Description
Current Time – Information on the system time and date on the device. If the current time has not been acquired by the SNTP client on the device or configured manually, this section shows the default time and date plus the time elapsed since the last system reset.	
<i>Time</i>	The current time on the system clock. This time is used to provide time stamps on log messages.
<i>Zone</i>	The acronym that represents the time zone.
<i>Date</i>	The current date on the system.
<i>Time Source</i>	The time source from which the time update is taken: SNTP The time has been acquired from an SNTP server. No Time Source The time has either been manually configured or not configured at all.
Time Zone – This section contains information about the time zone and offset.	
<i>Zone</i>	The acronym that represents the time zone.
<i>Offset</i>	The offset in hours from Coordinated Universal Time (UTC), also known as Greenwich Mean Time (GMT).
Summer Time – This section contains information on Summer Time (Daylight Saving Time).	
<i>Summer Time</i>	The summer time mode on the system: Disable Summer time is not active, and the time does not shift based on the time of year. Recurring Summer time occurs every year at the manually configured start and end dates and times. EU The system clock uses the standard recurring summer time settings for the European Union. All fields on the page except <i>Offset</i> and <i>Zone</i> are automatically populated and are not editable. USA The system clock uses the standard recurring daylight saving time settings for the United States. All fields on the page except <i>Offset</i> and <i>Zone</i> are automatically populated and are not editable. Non-Recurring Summer time settings are in effect only between the start date and end date of the specified year. If this mode is selected, the summer time settings do not repeat on an annual basis.
<i>Zone</i>	The acronym that represents the time zone of the summer time.

Time Zone Summary Fields (Continued)

Field	Description
<i>Offset</i>	The offset in hours from Coordinated Universal Time (UTC), also known as Greenwich Mean Time (GMT).
<i>Status</i>	Indicates if summer time is currently active.

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

Time Zone Configuration

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

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

*Time Zone Configuration**Time Zone Configuration Fields*

Field	Description
<i>Time Zone</i> – This section contains the following time zone settings:	
<i>Offset</i>	The system clock's offset from UTC, which is also known as Greenwich Mean Time (GMT).
<i>Zone</i>	The acronym that represents the time zone. This field is not validated against an official list of time zone acronyms.
<i>Date and Time</i> – Use the fields in this section to manually configure the system time and date. If the SNTP client is enabled (Unicast mode or Broadcast mode), these fields cannot be configured.	
<i>Time</i>	The current time in hours, minutes, and seconds on the system clock.
<i>Date</i>	The current date in month, day, and year on the system clock. To change the date, click  next to the field, select the year from the menu, browse to the desired month, and click the date.

Use the buttons to perform the following tasks:

- If you make any change to the page, click **Submit** to apply the settings.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Summer Time Configuration

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

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

The screenshot shows the 'Summer Time Configuration' page with three tabs: 'Summary', 'Time Zone', and 'Summer Time'. The 'Summer Time' tab is active. The page title is 'Summer Time Configuration'. Below the title, there is a 'Summer Time' dropdown menu set to 'Disable'. The page is divided into three main sections: 'Date Range', 'Recurring Date', and 'Zone'. Each section contains several input fields and dropdown menus. The 'Date Range' section includes 'Start Date', 'Starting Time of Day', 'End Date', and 'Ending Time of Day'. The 'Recurring Date' section includes 'Start Week', 'Start Day', 'Start Month', 'Starting Time of Day', 'End Week', 'End Day', 'End Month', and 'Ending Time of Day'. The 'Zone' section includes 'Offset' and 'Zone'. At the bottom of the page, there are three buttons: 'Submit', 'Refresh', and 'Cancel'. A copyright notice '© Copyright 2013-2014 Ubiquiti Networks, Inc.' is visible in the bottom right corner.

Section	Field	Value/Options
Date Range	Summer Time	Disable
	Start Date	[Calendar Icon]
	Starting Time of Day	[Time Picker] (00:00 to 23:59)
	End Date	[Calendar Icon]
Recurring Date	Ending Time of Day	[Time Picker] (00:00 to 23:59)
	Start Week	First
	Start Day	Sunday
	Start Month	January
	Starting Time of Day	[Time Picker] (00:00 to 23:59)
	End Week	First
	End Day	Sunday
	End Month	January
Zone	Ending Time of Day	[Time Picker] (00:00 to 23:59)
	Offset	[Time Picker] (-1 to 1440)
Zone	[Text Field]	(0 to 4 characters)

Buttons: Submit, Refresh, Cancel

© Copyright 2013-2014 Ubiquiti Networks, Inc.

Summer Time Configuration

Summer Time Configuration Fields

Field	Description
<i>Summer Time</i>	The summer time mode on the system: • Disable Summer time is not active, and the time does not shift based on the time of year. • Recurring Summer time occurs at the same time every year. The start and end times and dates for the time shift must be manually configured. • EU The system clock uses the standard recurring summer time settings used in countries in the European Union. When this field is selected, the rest of the applicable fields on the page except <i>Offset</i> and <i>Zone</i> are automatically populated and cannot be edited. • USA The system clock uses the standard recurring daylight saving time settings used in the United States. When this field is selected, the rest of the applicable fields on the page except <i>Offset</i> and <i>Zone</i> are automatically populated and cannot be edited. • Non-Recurring Summer time settings are in effect only between the start date and end date of the specified year. When this mode is selected, the summer time settings do not repeat on an annual basis.
<i>Date Range</i> – The fields in this section are available only if the <i>Summer Time</i> field is set to Non-Recurring mode.	
<i>Start Date</i>	The day, month, and year that summer time begins. To change the date, click  next to the field, select the year from the menu, browse to the desired month, and click the date.
<i>Starting Time and Day</i>	The time, in hours and minutes, to start summer time on the specified day.
<i>End Date</i>	The day, month, and year that summer time ends. To change the date, click  next to the field, select the year from the menu, browse to the desired month, and click the date.
<i>Ending Time of Day</i>	The time, in hours and minutes to end summer time on the specified day.
<i>Recurring Date</i> – The fields in this section are available only if the <i>Summer Time</i> field is set to Recurring mode.	
<i>Start Week</i>	The week of the month within which summer time begins.
<i>Start Day</i>	The day of the week on which summer time begins.
<i>Start Month</i>	The month of the year within which summer time begins.
<i>Starting Time of Day</i>	The time, in hours and minutes, to start summer time.
<i>End Week</i>	The week of the month within which summer time ends.
<i>End Day</i>	The day of the week on which summer time ends.
<i>End Month</i>	The month of the year within which summer time ends.
<i>Ending Time of Day</i>	The time, in hours and minutes, to end summer time.
<i>Zone</i> – The fields in this section are available for all modes selected from the <i>Summer Time</i> field except Disable .	
<i>Offset</i>	The number of minutes to shift the summer time from the standard time.
<i>Zone</i>	The acronym associated with the time zone when summer time is in effect.

Use the buttons to perform the following tasks:

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

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Chapter 4: Configuring Switching Information

- **[“Managing VLANs” on page 127](#)**
- **[“Creating MAC Filters” on page 134](#)**
- **[“GARP Configuration” on page 135](#)**
- **[“Configuring DHCP Snooping” on page 137](#)**
- **[“Configuring IGMP Snooping” on page 144](#)**
- **[“Configuring IGMP Snooping Querier” on page 151](#)**
- **[“Creating Port Channels” on page 154](#)**
- **[“Viewing Multicast Forwarding Database Information” on page 157](#)**
- **[“Configuring Protected Ports” on page 160](#)**
- **[“Configuring Spanning Tree Protocol” on page 161](#)**
- **[“Mapping 802.1p Priority” on page 170](#)**
- **[“Configuring Port Security” on page 172](#)**
- **[“Managing LLDP” on page 176](#)**

Managing VLANs

Adding Virtual LAN (VLAN) support to a Layer-2 switch offers some of the benefits of both bridging and routing. Like a bridge, a VLAN switch forwards traffic based on the Layer-2 header, which is fast, and like a router, it partitions the network into logical segments, which provides better administration, security and management of multicast traffic.

A VLAN is a set of end stations and the switch ports that connect them. You may have many reasons for the logical division, such as department or project membership. The only physical requirement is that the end station and the port to which it is connected both belong to the same VLAN.

Each VLAN in a network has an associated VLAN ID, which appears in the IEEE 802.1Q tag in the Layer-2 header of packets transmitted on a VLAN. An end station may omit the tag, or the VLAN portion of the tag, in which case the first switch port to receive the packet may either reject it or insert a tag using its default VLAN ID.

A given port may handle traffic for more than one VLAN, but it can only support one default VLAN ID.

VLAN Status

Use the *VLAN Status* page to view information about the VLANs configured on your system. To access the *VLAN Status* page, click **Switching > VLAN > Status** in the navigation menu.



VLAN Status

VLAN Status Fields

Field	Description
VLAN ID	The VLAN Identifier (VID) of the VLAN. The range of the VLAN ID is 1 to 4093. VLAN ID 1 is reserved for the default VLAN which is always present and cannot be edited or removed.
Name	The name of the VLAN. VLAN ID 1 is always named <i>default</i> .
Type	The VLAN type, which can be one of the following: Default The default VLAN which is always present Static A VLAN that you have created and configured Dynamic A VLAN created by GVRP registration that you have not converted to static, and that GVRP may therefore remove.
RSPAN	Displays <i>Enabled</i> if the VLAN is configured as the Remote Switched Port Analyzer (RSPAN) VLAN; otherwise, blank. The RSPAN VLAN is used to carry mirrored traffic from source ports to a destination probe port on a remote device.

Use the buttons to perform the following tasks:

- To add a VLAN, click **Add**. In the *Add VLAN* dialog box, specify the VLAN ID(s) in the *VLAN ID or Range* field (use "-" or "," to indicate a range of IDs or nonconsecutive IDs). Then, click **Submit** to apply the change.
- To configure a name for a VLAN or to convert a dynamic VLAN to a static VLAN, select the entry to modify and click **Edit**; then, configure the desired VLAN settings and click **Submit** to apply the changes.
- To remove one or more configured VLANs, select each entry to delete, click **Remove**, and confirm the deletion.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

VLAN Port Configuration

Use the *VLAN Port Configuration* page to configure a virtual LAN on a port.

To access the page, click **Switching > VLAN > Port Configuration** in the navigation menu.

VLAN Port Configuration

VLAN Port Configuration Fields

Field	Description
<i>VLAN ID</i>	The menu includes the VLAN ID for all VLANs configured on the device. To view or configure settings for a VLAN, be sure to select the correct VLAN from the menu.
<i>Interface</i>	The interface associated with the rest of the data in the row. When editing VLAN information for one or more interfaces, this field identifies the interfaces that are being configured.
<i>Status</i>	The current participation mode of the interface in the selected VLAN. The <i>Status</i> value differs from the <i>Participation</i> value only when the <i>Participation</i> mode is <i>Auto Detect</i> . The <i>Status</i> is one of the following: Include The port is a member of the selected VLAN. Exclude The port is not a member of the selected VLAN.
<i>Participation</i>	The participation mode of the interface in the selected VLAN, which is one of the following: Include The port is always a member of the selected VLAN. This mode is equivalent to registration fixed in the IEEE 802.1Q standard. Exclude The port is never a member of the selected VLAN. This mode is equivalent to registration forbidden in the IEEE 802.1Q standard. Auto Detect The port can be dynamically registered in the selected VLAN through GVRP. The port will not participate in this VLAN unless it receives a GVRP request. This mode is equivalent to registration normal in the IEEE 802.1Q standard.
<i>Tagging</i>	The tagging behavior for all the ports in this VLAN, which is one of the following: Tagged The frames transmitted in this VLAN will include a VLAN ID tag in the Ethernet header. Untagged The frames transmitted in this VLAN will be untagged.

Use the buttons to perform the following tasks:

- To configure a specific interface in the selected VLAN, select the interface's entry from table, click **Edit**, configure the fields as needed, and click **Submit** to apply the changes.
- To configure all interfaces in the selected VLAN, click **Edit All**, configure the fields as needed, and click **Submit** to apply the changes.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

VLAN Port Summary

Use the *VLAN Port Summary* page to view VLAN configuration information for all the ports on the system. To access the *VLAN Port Summary* page, click **Switching > VLAN > Port Summary** in the navigation menu.

Interface	Port VLAN ID	Acceptable Frame Type	Ingress Filtering	Priority
0/1	1	Admit All	Disable	0
0/2	1	Admit All	Disable	0
0/3	1	Admit All	Disable	0
0/4	1	Admit All	Disable	0
0/5	1	Admit All	Disable	0
0/6	1	Admit All	Disable	0
0/7	1	Admit All	Disable	0
0/8	1	Admit All	Disable	0
0/9	1	Admit All	Disable	0
0/10	1	Admit All	Disable	0

VLAN Port Summary

VLAN Port Summary Fields

Field	Description
<i>Interface</i>	Identifies the physical interface associated with the rest of the data in the row.
<i>Port VLAN ID</i>	The VLAN ID assigned to untagged or priority tagged frames received on this port. This value is also known as the Port VLAN ID (PVID). In a tagged frame, the VLAN is identified by the VLAN ID in the tag.
<i>Acceptable Frame Types</i>	Indicates how the interface handles untagged and priority tagged frames. The options include: Admit All Untagged and priority tagged frames received on the interface are accepted and assigned the value of the Port VLAN ID for this interface. Only Tagged The interface discards any untagged or priority tagged frames it receives. Only Untagged The interface discards any tagged frames it receives. For all options, VLAN tagged frames are forwarded in accordance with the IEEE 802.1Q VLAN standard.
<i>Ingress Filtering</i>	Shows how the port handles tagged frames. Enable Discards a frame if the VLAN ID in the tag identifies a VLAN to which the port does not belong. Disable Accepts all tagged frames (factory default).
<i>Priority</i>	Identifies the default 802.1p priority assigned to untagged packets arriving at the port.

Use the buttons to perform the following tasks:

- To configure a specific port interface, select the interface's entry from table, click **Edit**, configure the fields as needed, and click **Submit** to apply the changes.
- To configure all port interfaces, click **Edit All**, configure the fields as needed, and click **Submit** to apply the changes.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

VLAN Internal Usage

Use the *VLAN Internal Usage Configuration* page to assign a Base VLAN ID for internal allocation of VLANs to the routing interface.

To access the *VLAN Internal Usage Configuration* page, click **Switching > VLAN > Internal Usage** in the navigation menu.

VLAN Internal Usage

VLAN Internal Usage Fields

Field	Description
<i>Base VLAN ID</i>	The first VLAN ID to be assigned to a port-based routing interface.
<i>Allocation Policy</i>	Determines whether VLAN IDs assigned to port-based routing interfaces start at the base and decrease in value (<i>Descending</i>) or start at the base and increase in value (<i>Ascending</i>).
<i>VLAN ID</i>	The VLAN ID assigned to a port-based routing interface. The device automatically assigns an unused VLAN ID when the routing interface is created.
<i>Routing Interface</i>	The port-based routing interface associated with the VLAN.

Use the buttons to perform the following tasks:

- If you change any information on the page, click **Submit** to apply the changes to the system.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Reset VLAN Configuration

Use the *Reset VLAN Configuration* page to return all VLAN parameters for all interfaces to the factory default values.

To access the *Reset VLAN Configuration* page, click **Switching** > **VLAN** > **Reset** in the navigation menu.



Reset VLAN Configuration

To reset the VLAN configuration, click **Reset**, and then confirm the reset by clicking **OK**. When the system indicates that all default VLAN settings have been restored, click **Close** to acknowledge the result.

Managing Voice VLANs

Voice VLAN Configuration

The voice VLAN feature enables switch ports to carry voice traffic with defined settings so that voice and data traffic are separated when coming onto the port. A voice VLAN ensures that the sound quality of an IP phone is safeguarded from deterioration when data traffic on the port is high.

The inherent isolation provided by VLANs ensures that inter-VLAN traffic is under management control and that network-attached clients cannot initiate a direct attack on voice components. A QoS protocol based on the IEEE 802.1P class-of-service (CoS) protocol uses classification and scheduling to send network traffic from the switch in a predictable manner. The system uses the source MAC of the traffic traveling through the port to identify the IP phone data flow.

Voice VLAN is enabled per-port basis. A port can participate only in one voice VLAN at a time. The Voice VLAN feature is disabled by default.

To display the *Voice VLAN Configuration* page, click **Switching > Voice VLAN > Configuration**.



Voice VLAN Configuration

Voice VLAN Configuration Fields

Field	Description
<i>Voice VLAN Admin Mode</i>	The administrative mode of the Voice VLAN feature. Click <i>Enable</i> or <i>Disable</i> (default) to administratively turn the Voice VLAN feature on or off for all ports. When Voice VLAN is enabled globally and configured on interfaces that carry voice traffic, this feature can help ensure that the sound quality of an IP phone does not deteriorate when data traffic on the port is high.

Use the buttons to perform the following tasks:

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

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Voice VLAN Interface Summary

Use this page to configure the per-port settings for the Voice VLAN feature. When Voice VLAN is configured on a port that receives both voice and data traffic, it can help ensure that the voice traffic has priority.

To display the *Voice VLAN Interface Summary* page, click **Switching > Voice VLAN > Interface Summary**.



Voice VLAN Interface Summary

Voice VLAN Interface Summary Fields

Field	Description
<i>Interface</i>	The interface associated with the rest of the data in the row. When adding a Voice VLAN configuration to a port, the Interface menu allows you to select the port to configure. Only interfaces that have not been configured with Voice VLAN settings can be selected from the menu.
<i>Operational State</i>	The operational status of the Voice VLAN feature on the interface. To be enabled, Voice VLAN must be globally enabled and enabled on the interface. Additionally, the interface must be up and have a link.
<i>CoS Override Mode</i>	The Class of Service override mode: Enabled The port ignores the 802.1p priority value in the Ethernet frames it receives from connected devices. Disabled The port trusts the priority value in the received frame.
<i>Voice VLAN Interface Mode</i>	Indicates how an IP phone connected to the port should send voice traffic: VLAN ID Forward voice traffic in the specified voice VLAN. Dot1p Tag voice traffic with the specified 802.1p priority value. None Use the settings configured on the IP phone to send untagged voice traffic. Untagged Send untagged voice traffic. Disable Operationally disables the Voice VLAN feature on the interface.
<i>Add Voice VLAN and Edit Voice VLAN dialog boxes</i> – When you click Add or Edit , the following configurable field is displayed:	
<i>Voice VLAN Interface Value</i>	When adding or editing Voice VLAN settings for an interface and either VLAN ID or Dot1p is selected as the <i>Voice VLAN Interface Mode</i> , specify the voice VLAN ID or the Dot1p priority value that the connected IP phone should use for voice traffic.

Use the buttons to perform the following tasks:

- To configure Voice VLAN settings on a port, click **Add**. Select the interface to configure from the *Interface* drop-down menu, configure the remaining settings, and then click **Submit** to apply the changes.
- To change the Voice VLAN settings, select the interface to modify and click **Edit**, configure the settings as needed, and then click **Submit** to apply the changes.
- To remove the Voice VLAN configuration from one or more ports, select each entry to delete, click **Remove**, and confirm the deletion.
- Click **Refresh** to refresh the page with the most current data from the switch.


To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Creating MAC Filters

Static MAC filtering allows you to associate a MAC address with a VLAN and set of source ports and destination ports. (The availability of source and destination port filters is subject to platform restrictions). Any packet with a static MAC address in a specific VLAN is admitted only if the ingress port is included in the set of source ports; otherwise the packet is dropped. If admitted, the packet is forwarded to all the ports in the destination list.

MAC Filter Configuration

Use the *MAC Filter Configuration* page to associate a MAC address with a VLAN and one or more source and/or destination ports. To access the page, click **Switching** > **Filters** > **MAC Filters** in the navigation menu.



MAC Filter Configuration

MAC Filter Configuration Fields

Field	Description
<i>MAC Address</i>	The MAC address of the filter. The destination MAC address of an Ethernet frame must match this value to be considered for the filter. When adding or editing a filter, note that you cannot configure the following MAC addresses in this field: • 00:00:00:00:00:00 • 01:80:C2:00:00:00 to 01:80:C2:00:00:0F • 01:80:C2:00:00:20 to 01:80:C2:00:00:21 • FF:FF:FF:FF:FF:FF
<i>VLAN ID</i>	The VLAN ID associated with the filter. The VLAN ID is used with the MAC address to fully identify the frames to filter.
<i>Source Members</i>	The port(s) included in the inbound filter. If a frame with the MAC address and VLAN ID specified by the filter arrives on a port in the <i>Source Members</i> list, it is forwarded to a port in the <i>Destination Members</i> list. If the frame that meets the filter criteria arrives on a port that is not in the <i>Source Members</i> list, it is dropped.
<i>Destination Members</i>	The port(s) included in the outbound filter. A frame with the MAC address and VLAN ID combination specified in the filter is transmitted only out of ports in the list.

Use the buttons to perform the following tasks:

- To add a MAC filter, click **Add**. In the *Add Static MAC Filter* dialog box, enter a valid *MAC Address*, select a *VLAN ID* (the drop-down box only lists VLANs currently configured on the system), select the desired source and destination port(s) from the *Available Port List* fields (press and hold CTRL to select multiple ports), and click to move the selected ports to the *Source Members* and *Destination Members* fields. Then, click **Submit** to apply the changes.
- To change the source and/or destination members for an existing filter, select the filter's entry from the table, and click **Edit**. When you have completed the changes, click **Submit** to apply the changes.
- To remove a filter, select it from the table, click **Remove**, and confirm the deletion.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

GARP Configuration

Use this page to set the administrative mode for the features that use the Generic Attribute Registration Protocol (GARP), including GARP VLAN Registration Protocol (GVRP) and GARP Multicast Registration Protocol (GMRP). GARP is a general-purpose protocol that registers any network connectivity or membership-style information. GARP defines a set of switches interested in a given network attribute, such as VLAN ID or multicast address.

GARP Switch Configuration

To access the *GARP Switch Configuration* page, click **Switching** > **GARP** > **Switch** in the navigation menu.

GARP Switch Configuration

GARP Switch Configuration Fields

Field	Description
<i>GVRP Mode</i>	The administrative mode of GVRP on the system. When set to <i>Enable</i> , GVRP can help dynamically manage VLAN memberships on trunk ports.
<i>GMRP Mode</i>	The administrative mode of GMRP on the system. When set to <i>Enable</i> , GMRP can help control the flooding of multicast traffic by keeping track of group membership information. GMRP is similar to IGMP snooping in its purpose, but IGMP snooping is more widely used. GMRP must be running on both the host and the switch to function properly.

Use the buttons to perform the following tasks:

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

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

GARP Port Configuration

Use this page to set the per-interface administrative mode for GARP VLAN Registration Protocol (GVRP) and GARP Multicast Registration Protocol (GMRP). On this page, you can also set the GARP timers for each interface. GVRP and GMRP use the same set of GARP timers to specify the amount of time to wait before transmitting various GARP messages.

To access the *GARP Port Configuration* page, click **Switching** > **GARP** > **Port** in the navigation menu.

Interface	GVRP Mode	GMRP Mode	Join Timer (Centisecs)	Leave Timer (Centisecs)	Leave All Timer (Centisecs)
0/1	Disabled	Disabled	20	60	1000
0/2	Disabled	Disabled	20	60	1000
0/3	Disabled	Disabled	20	60	1000
0/4	Disabled	Disabled	20	60	1000
0/5	Disabled	Disabled	20	60	1000
0/6	Disabled	Disabled	20	60	1000
0/7	Disabled	Disabled	20	60	1000
0/8	Disabled	Disabled	20	60	1000
0/9	Disabled	Disabled	20	60	1000
0/10	Disabled	Disabled	20	60	1000

GARP Port Configuration

GARP Port Configuration Fields

Field	Description
<i>Interface</i>	The interface associated with the rest of the data in the row. When configuring one or more interfaces in the Edit GARP Port Configuration window, this field identifies the interfaces that are being configured.
<i>GVRP Mode</i>	The administrative mode of GVRP on the interface. When enabled, GVRP can help dynamically manage VLAN memberships on trunk ports. GVRP must also be enabled globally for the protocol to be active on the interface. When disabled, the protocol will not be active on the interface, and the GARP timers have no effect.
<i>GMRP Mode</i>	The administrative mode of GMRP on the interface. When enabled, GMRP can help control the flooding of multicast traffic by keeping track of group membership information. GMRP must also be enabled globally for the protocol to be active on the interface. When disabled, the protocol will not be active on the interface, and the GARP timers have no effect.
<i>Join Timer (Centisecs)</i>	The amount of time between the transmission of GARP PDUs registering (or re-registering) membership for a VLAN or multicast group.
<i>Leave Timer (Centisecs)</i>	The amount of time to wait after receiving an unregister request for a VLAN or multicast group before deleting the associated entry. This timer allows time for another station to assert registration for the same attribute in order to maintain uninterrupted service.
<i>Leave All Timer (Centisecs)</i>	The amount of time to wait before sending a LeaveAll PDU after the GARP application has been enabled on the interface or the last LeaveAll PDU was sent. A LeaveAll PDU indicates that all registrations will shortly be deregistered. Participants will need to rejoin in order to maintain registration.

To change the GARP settings for one or more interfaces, select each interface to configure and click **Edit**. The same settings are applied to all selected interfaces.

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

Configuring DHCP Snooping

DHCP snooping is a security feature that monitors DHCP messages between a DHCP client and DHCP servers to filter harmful DHCP messages and to build a bindings database of {MAC address, IP address, VLAN ID, port} tuples that are considered authorized. You can enable DHCP snooping globally and on specific VLANs, and configure ports within the VLAN to be trusted or untrusted. If a DHCP message arrives on an untrusted port, DHCP snooping filters messages that are not from authorized DHCP clients. DHCP server messages are forwarded only through trusted ports.

Global DHCP Snooping Configuration

Use this page to view and configure the global settings for DHCP Snooping.

To access the *Global DHCP Snooping Configuration* page, click **Switching > DHCP Snooping > Base > Global** in the navigation menu.

Global DHCP Snooping Configuration

Global DHCP Snooping Configuration Fields

Field	Description
<i>DHCP Snooping Mode</i>	Used to <i>Enable</i> or <i>Disable</i> DHCP snooping on the device.
<i>MAC Address Validation</i>	Used to <i>Enable</i> or <i>Disable</i> the verification of the sender MAC address for DHCP snooping. When enabled, the device checks packets that are received on untrusted interface to verify that the MAC address and the DHCP client hardware address match. If the addresses do not match, the device drops the packet.

Use the buttons to perform the following tasks:

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

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

DHCP Snooping VLAN Configuration

Use this page to view and configure the DHCP snooping settings on VLANs that exist on the device. DHCP snooping can be configured on switching VLANs and routing VLANs. For Layer-2 (non-routing) VLANs, DHCP snooping forwards valid DHCP client messages received on the VLANs. The message is forwarded on all trusted interfaces in the VLAN. When a DHCP packet is received on a routing VLAN, the DHCP snooping application applies its filtering rules and updates the bindings database. If a client message passes filtering rules, the message is placed into the software forwarding path, where it may be processed by the DHCP relay agent, the local DHCP server, or forwarded as an IP packet.

To access the *DHCP Snooping VLAN Configuration* page, click **Switching > DHCP Snooping > Base > VLAN Configuration** in the navigation menu.



DHCP Snooping VLAN Configuration

DHCP Snooping VLAN Configuration Fields

Field	Description
<i>VLAN ID</i>	The VLAN ID that is enabled for DHCP snooping. In the Add DHCP Snooping VLAN Configuration window, this field lists the VLAN ID of all VLANs that exist on the device.
<i>DHCP Snooping Mode</i>	The current administration mode (<i>Enabled</i> or <i>Disabled</i>) of DHCP snooping for the VLAN. Only VLANs that are enabled for DHCP snooping appear in the list.

Use the buttons to perform the following tasks:

- To enable a VLAN for DHCP snooping, click **Add** and select the VLAN (press and hold CTRL to select multiple VLANs). Then, click **Submit** to apply the changes.
- To disable DHCP snooping on one or more VLANs, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

DHCP Snooping Interface Configuration

Use this page to view and configure the DHCP snooping settings for each interface. The DHCP snooping feature processes incoming DHCP messages. For DHCPRELEASE and DHCPDECLINE messages, the feature compares the receive interface and VLAN with the client's interface and VLAN in the binding database. If the interfaces do not match, the application logs the event (when logging of invalid packets is enabled) and drops the message. If MAC address validation is globally enabled, messages that pass the initial validation are checked to verify that the source MAC address and the DHCP client hardware address match. Where there is a mismatch, DHCP snooping logs the event (when logging of invalid packets is enabled) and drops the packet. To change the DHCP Snooping settings for one or more interfaces, select each entry to modify and click **Edit**. The same settings are applied to all selected interfaces.

To access the *DHCP Snooping Interface Configuration* page, click **Switching > DHCP Snooping > Base > Interface Configuration** in the navigation menu.

The screenshot displays the 'DHCP Snooping Interface Configuration' page. At the top, there are tabs for 'Global', 'VLAN Configuration', 'Interface Configuration' (selected), 'Static Bindings', 'Dynamic Bindings', 'Persistent', and 'Statistics'. Below the tabs, the title 'DHCP Snooping Interface Configuration' is shown. A 'Display' dropdown is set to '10' rows, and it indicates 'Showing 1 to 10 of 32 entries'. A 'Filter' input field is present. The main table has the following columns: 'Interface', 'Trust State', 'Log Invalid Packets', 'Rate Limit (pps)', and 'Burst Interval (Seconds)'. The table lists interfaces Q/1 through Q/10, all with 'Disabled' for both 'Trust State' and 'Log Invalid Packets'. At the bottom of the table, there are navigation buttons: 'First', 'Previous', '1', '2', '3', '4', 'Next', 'Last', 'Refresh', and 'Edit'.

DHCP Snooping Interface Configuration

DHCP Snooping Interface Configuration Fields

Field	Description
<i>Interface</i>	The interface associated with the rest of the data in the row. When configuring the settings for one or more interfaces, this field identifies each interface that is being configured.
<i>Trust State</i>	The trust state configured on the interface. The trust state is one of the following: Disabled The interface is considered to be untrusted and could potentially be used to launch a network attack. DHCP server messages are checked against the bindings database. On untrusted ports, DHCP snooping enforces the following security rules: - DHCP packets from a DHCP server (DHCP OFFER, DHCP ACK, DHCP NAK, DHCP RELEASE QUERY) are dropped. - DHCP RELEASE and DHCP DECLINE messages are dropped if the MAC address is in the snooping database but the binding's interface is other than the interface where the message was received. - DHCP packets are dropped when the source MAC address does not match the client hardware address if MAC Address Validation is globally enabled. Enabled The interface is considered trusted and forwards DHCP server messages without validation.
<i>Log Invalid Packets</i>	The administrative mode of invalid packet logging on the interface. If enabled, the DHCP snooping feature generates a log message when an invalid packet is received and dropped by the interface.
<i>Rate Limit (pps)</i>	The rate limit value for DHCP packets received on the interface. To prevent DHCP packets from being used as a DoS attack when DHCP snooping is enabled, the snooping application enforces a rate limit for DHCP packets received on untrusted interfaces. If the incoming rate of DHCP packets exceeds the value of this object during the amount of time specified for the burst interval, the port will be shut down. You must administratively enable the port to allow it to resume traffic forwarding.
<i>Burst Interval (Seconds)</i>	The burst interval value for rate limiting on this interface. If the rate limit is unspecified, then burst interval has no meaning.

Use the buttons to perform the following tasks:

- To edit DHCP snooping on one or more interfaces, select each interface entry, click **Edit**, change the settings as needed, and click **Submit** to apply the changes.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

DHCP Snooping Static Bindings

Use this page to view, add, and remove static bindings in the DHCP snooping bindings database.

To access the *DHCP Snooping Static Bindings* page, click **Switching** > **DHCP Snooping** > **Base** > **Static Bindings** in the navigation menu.



DHCP Snooping Static Bindings

DHCP Snooping Static Bindings Fields

Field	Description
<i>Interface</i>	The interface on which the DHCP client is authorized.
<i>MAC Address</i>	The MAC address associated with the DHCP client. This is the Key to the binding database.
<i>VLAN ID</i>	The ID of the VLAN the client is authorized to use.
<i>IP Address</i>	The IP address of the client.

Use the buttons to perform the following tasks:


- To add a DHCP Snooping static binding, click **Add**. Then, enter a valid MAC address, select a VLAN ID from the drop-down menu, enter an IP address, and click **Submit** to apply the changes.
- To remove a DHCP snooping static binding, select the entry from the table, click **Remove**, and confirm the operation.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

DHCP Snooping Dynamic Bindings

Use this page to view and clear dynamic bindings in the DHCP snooping bindings database. The DHCP snooping feature uses DHCP messages to build and maintain the bindings database. The bindings database includes data for clients only on untrusted ports. DHCP snooping creates a tentative binding from DHCP DISCOVER and REQUEST messages. Tentative bindings tie a client to an interface (the interface where the DHCP client message was received). Tentative bindings are completed when DHCP snooping learns the client's IP address from a DHCP ACK message on a trusted port. DHCP snooping removes bindings in response to DECLINE, RELEASE, and NACK messages. The DHCP snooping feature ignores the ACK messages as a reply to the DHCP Inform messages received on trusted ports.

To access the *DHCP Snooping Dynamic Bindings* page, click **Switching > DHCP Snooping > Base > Dynamic Bindings** in the navigation menu.



DHCP Snooping Dynamic Bindings

DHCP Snooping Dynamic Bindings Fields

Field	Description
<i>Interface</i>	The interface on which the DHCP client message was received.
<i>MAC Address</i>	The MAC address associated with the DHCP client that sent the message. This is the key to the binding database.
<i>VLAN ID</i>	The VLAN ID of the client interface.
<i>IP Address</i>	The IP address assigned to the client by the DHCP server.
<i>Lease Time</i>	The remaining IP address lease time for the client.

Use the buttons to perform the following tasks:

- To remove one or more entries in the database, select each entry to delete and click **Clear**. You must confirm the action before the entry is deleted.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

DHCP Snooping Persistent Configuration

Use this page to configure the persistent location of the DHCP snooping bindings database. The bindings database can be stored locally on the device or on a remote system somewhere else in the network. The device must be able to reach the IP address of the remote system to send bindings to a remote database.

To access the *DHCP Snooping Persistent Configuration* page, click **Switching > DHCP Snooping > Base > Persistent** in the navigation menu.

DHCP Snooping Persistent Configuration

DHCP Snooping Persistent Configuration Fields

Field	Description
<i>Store</i>	The location of the DHCP snooping bindings database, which is either locally on the device (<i>Local</i>) or on a remote system (<i>Remote</i>).
<i>Remote IP Address</i>	The IP address of the system on which the DHCP snooping bindings database will be stored. This field is available only if <i>Remote</i> is selected in the <i>Store</i> field.
<i>Remote File Name</i>	The file name of the DHCP snooping bindings database in which the bindings are stored. This field is available only if <i>Remote</i> is selected in the <i>Store</i> field.
<i>Write Delay (Seconds)</i>	The amount of time to wait between writing bindings information to persistent storage. This allows the device to collect as many entries as possible (new and removed) before writing them to the persistent file.

Use the buttons to perform the following tasks:

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

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

DHCP Snooping Statistics

Use this page to configure the persistent location of the DHCP snooping bindings database. The bindings database can be stored locally on the device or on a remote system somewhere else in the network. The device must be able to reach the IP address of the remote system to send bindings to a remote database.

To access the *DHCP Snooping Statistics* page, click **Switching > DHCP Snooping > Base > Statistics** in the navigation menu.

Interface	MAC Verify Failures	Client Ifc Mismatch	DHCP Server Msgs Received
0/1	0	0	0
0/2	0	0	0
0/3	0	0	0
0/4	0	0	0
0/5	0	0	0
0/6	0	0	0
0/7	0	0	0
0/8	0	0	0
0/9	0	0	0
0/10	0	0	0

DHCP Snooping Statistics

DHCP Snooping Statistics Fields

Field	Description
<i>Interface</i>	The interface associated with the rest of the data in the row.
<i>MAC Verify Failures</i>	The number of DHCP messages that were dropped because the source MAC address and client hardware address did not match. MAC address verification is performed only if it is globally enabled.
<i>Client Ifc Mismatch</i>	The number of packets that were dropped by DHCP snooping because the interface and VLAN on which the packet was received does not match the client's interface and VLAN information stored in the binding database.
<i>DHCP Server Msgs Received</i>	The number of DHCP server messages (DHCP OFFER, DHCP ACK, DHCP NAK, DHCP RELEASE QUERY) that have been dropped on an untrusted port.

Use the buttons to perform the following tasks:

- To reset the statistics to zero for one or more interfaces, select each interface with the data to reset and click **Clear Counters**. You must confirm the action before the counters are reset.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Configuring IGMP Snooping

Internet Group Management Protocol (IGMP) Snooping is a feature that allows a switch to forward multicast traffic intelligently on the switch. Multicast IP traffic is traffic that is destined to a host group. Host groups are identified by class D IP addresses, which range from 224.0.0.0 to 239.255.255.255. Based on the IGMP query and report messages, the switch forwards traffic only to the ports that request the multicast traffic. This prevents the switch from broadcasting the traffic to all ports and possibly affecting network performance.

A traditional Ethernet network may be separated into different network segments to prevent placing too many devices onto the same shared media. Bridges and switches connect these segments. When a packet with a broadcast or multicast destination address is received, the switch will forward a copy into each of the remaining network segments in accordance with the IEEE MAC Bridge standard. Eventually, the packet is made accessible to all nodes connected to the network.

This approach works well for broadcast packets that are intended to be seen or processed by all connected nodes. In the case of multicast packets, however, this approach could lead to less efficient use of network bandwidth, especially if the packets are for only a small number of nodes. Packets are flooded into network segments where no node has any interest in receiving them. While nodes rarely incur any processing overhead to filter packets addressed to unrequested group addresses, they are unable to transmit new packets onto the shared media for the period of time that the multicast packet is flooded. The problem of wasting bandwidth is even worse when the LAN segment is not shared, for example in Full Duplex links.

Allowing switches to snoop IGMP packets is a creative effort to solve this problem. The switch uses the information in the IGMP packets as they are being forwarded throughout the network to determine which segments should receive packets directed to the group address.

Global Configuration and Status

Use the *IGMP Snooping Global Configuration and Status* page to enable IGMP snooping on the switch and view information about the current IGMP configuration. To access the page, click **Switching > IGMP Snooping > Configuration and Status** in the navigation menu.

IGMP Snooping Global Configuration and Status

IGMP Snooping Global Configuration and Status Fields

Field	Description
<i>Admin Mode</i>	Select the administrative mode for IGMP Snooping for the switch. The default is <i>Disable</i> .
<i>Multicast Control Frame Count</i>	Shows the number of multicast control frames that have been processed by the CPU.
<i>Interfaces Enabled for IGMP Snooping</i>	Lists the interfaces currently enabled for IGMP Snooping. To enable interfaces for IGMP snooping, see "Interface Configuration" on page 145 .
<i>Data Frames Forwarded by CPU</i>	Shows the number of multicast control frames processed by the CPU.

Use the buttons to perform the following tasks:

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

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Interface Configuration

Use the *IGMP Snooping Interface Configuration* page to configure IGMP snooping settings on specific interfaces.

To access the page, click **Switching** > **IGMP Snooping** > **Interface Configuration** in the navigation menu.

IGMP Snooping Interface Configuration

IGMP Snooping Interface Configuration Fields

Field	Description
<i>Interface</i>	The physical or LAG interface.
<i>Admin Mode</i>	The interface mode for the selected interface for IGMP Snooping for the switch. The default is <i>Disable</i> . IGMP snooping must be enabled globally and on an interface for the interface to be able to snoop IGMP packets to determine which segments should receive multicast packets directed to the group address.
<i>Group Membership Interval</i>	The amount of time in seconds that the interface should wait for a report for a particular group on a particular interface before IGMP snooping deletes that interface from the group. The valid range is from 2 to 3600 seconds. The default is 260 seconds.
<i>Max Response Time</i>	The amount of time in seconds that the interface should wait after sending a query if it does not receive a report for a particular group on that interface. The value must be greater or equal to 1 and less than the <i>Group Membership Interval</i> in seconds. The default is 10 seconds.
<i>Multicast Router Expiration Time</i>	The amount of time in seconds that the interface should wait to receive a query on an interface before it is removed from the list of interfaces with multicast routers attached. The valid range is from 0 to 3600 seconds. The default is 0 seconds (indicates an infinite timeout; i.e., no expiration).
<i>Fast Leave Admin Mode</i>	The Fast Leave mode (default <i>Disable</i>) for an interface. If enabled, the interface can be immediately removed from the Layer-2 forwarding table entry upon receiving an IGMP leave message for a multicast group without first sending out MAC-based general queries.

Use the buttons to perform the following tasks:

- To edit the IGMP snooping configuration of one or more interfaces, select the interface(s), click **Edit**, and make the changes as needed. Then, click **Submit** to apply the changes.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

IGMP Snooping Source Specific Multicast

This page displays information about multicast groups discovered by snooping IGMPv3 reports.

To access the *IGMP Snooping Source Specific Multicast* page, click **Switching** > **IGMP Snooping** > **Source Specific Multicast** in the navigation menu.



IGMP Snooping Source Specific Multicast

IGMP Snooping Source Specific Multicast Fields


Field	Description
<i>VLAN ID</i>	VLAN on which the IGMP v3 report is received.
<i>Group</i>	The IPv4 multicast group address.
<i>Interface</i>	The interface on which the IGMP v3 report is received.
<i>Reporter</i>	The IPv4 address of the host that sent the IGMPv3 report.
<i>Source Filter Mode</i>	The source filter mode (<i>Include</i> or <i>Exclude</i>) for the specified group.
<i>Source Address List</i>	List of source IP addresses for which source filtering is requested.

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

IGMP Snooping VLAN Status

Use this page to enable or disable IGMP snooping on system VLANs and to view and configure per-VLAN IGMP snooping settings. Only VLANs that are enabled for IGMP snooping appear in the table.

To access the *IGMP Snooping VLAN Status* page, click **Switching > IGMP Snooping > VLAN Status** in the navigation menu.



IGMP Snooping VLAN Status

IGMP Snooping VLAN Status Fields

Field	Description
VLAN ID	The VLAN associated with the rest of the data in the row. When enabling IGMP snooping on a VLAN, use this menu to select the desired VLAN. Only VLANs that have been configured on the system and are not already enabled for IGMP snooping appear in the menu. When modifying IGMP snooping settings, this field identifies the VLAN that is being configured.
Admin Mode	The administrative mode of IGMP snooping on the VLAN. IGMP snooping must be enabled globally and on an VLAN for the VLAN to be able to snoop IGMP packets to determine which network segments should receive multicast packets directed to the group address.
Fast Leave Admin Mode	The administrative mode of Fast Leave on the VLAN. If Fast Leave is enabled, the VLAN can be immediately removed from the Layer-2 forwarding table entry upon receiving an IGMP leave message for a multicast group without first sending out MAC-based general queries.
Group Membership Interval (Seconds)	The number of seconds the VLAN should wait for a report for a particular group on the VLAN before the IGMP snooping feature deletes the VLAN from the group.
Max Response Time (Seconds)	The number of seconds the VLAN should wait after sending a query if does not receive a report for a particular group. The specified value should be less than the <i>Group Membership Interval</i> .
Multicast Router Expiration Time (Seconds)	The number of seconds the VLAN should wait to receive a query before it is removed from the list of VLANs with multicast routers attached.
Report Suppression Mode	The IGMPv1 and IGMPv2 report suppression mode. The device uses IGMP report suppression to limit the membership report traffic sent to multicast-capable routers. When this mode is enabled, the device does not send duplicate reports to the multicast router. Note that this mode is supported only when the multicast query has IGMPv1 and IGMPv2 reports. This feature is not supported when the query includes IGMPv3 reports. The options are as follows: Enabled Only the first IGMP report from all hosts for a group IGMP report is forwarded to the multicast routers. Disabled The device forwards all IGMP reports from all hosts in a multicast group to the multicast routers.

Use the buttons to perform the following tasks:


- To enable IGMP snooping on a VLAN, click **Add**, configure the settings in the fields, and click **Submit** to apply the changes.
- To change the IGMP snooping settings of an IGMP snooping-enabled VLAN, select the entry to be changed, click **Edit**, and make the changes as needed. Then, click **Submit** to apply the changes.
- To disable IGMP snooping on one or more VLANs, select each VLAN, click **Remove**, and confirm the action. When IGMP snooping is disabled, the VLAN entry is removed from the table, but the VLAN itself still exists on the system.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

IGMP Snooping Multicast Router Configuration

If a multicast router is attached to the switch, its existence can be learned dynamically. You can also statically configure a switch port as a multicast router interface. Use the *IGMP Snooping Multicast Router Configuration* page to manually configure an interface as a static multicast router interface.

To access the *IGMP Snooping Multicast Router Configuration* page, click **Switching** > **IGMP Snooping** > **Multicast Router Configuration** in the navigation menu.



IGMP Snooping Multicast Router Configuration

IGMP Snooping Multicast Router Configuration Fields

Field	Description
<i>Interface</i>	Select the physical or LAG interface to display.
<i>Multicast Router</i>	Set the multicast router status: Enabled The port is a multicast router interface. Disabled The port does not have a multicast router configured.

Use the buttons to perform the following tasks:

- To or disable multicast router configuration on one or more interfaces, select the interfaces, click **Edit**, configure the *Multicast Router* field, and click **Submit** to apply the changes.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

IGMP Snooping Multicast Router VLAN Status

If a multicast router is attached to the switch, its existence can be learned dynamically. You can also statically configure one or more VLANs on each interface to act as a multicast router interface, which is an interface that faces a multicast router or IGMP querier and receives multicast traffic.

To access the *IGMP Snooping Multicast Router VLAN Status* page, click **Switching > IGMP Snooping > Multicast Router VLAN Status** in the navigation menu.



IGMP Snooping Multicast Router VLAN Status

IGMP Snooping Multicast Router VLAN Status Fields

Field	Description
<i>Interface</i>	The interface associated with the rest of the data in the row. Only interfaces that are configured with multicast router VLANs appear in the table.
<i>VLAN IDs</i>	The ID of the VLAN configured as enabled for multicast routing on the associated interface.

Use this page to view the multicast router VLAN status for each interface. Use the buttons to perform the following tasks:

- Click the **Add** and **Edit** buttons to be redirected to the *Multicast Router VLAN Configuration* page for the selected interface to enable or disable VLANs as multicast router interfaces.
- To disable all VLANs as multicast router interfaces for one or more physical ports or LAGs, select each entry to modify, click **Remove**, and confirm the action.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

IGMP Snooping Multicast Router VLAN Configuration

Use this page to enable or disable specific VLANs as multicast router interfaces for a physical port or LAG. A multicast router interface faces a multicast router or IGMP querier and receives multicast traffic.

To access the *IGMP Snooping Multicast Router VLAN Configuration* page, click **Switching > IGMP Snooping > Multicast Router VLAN Configuration** in the navigation menu.

IGMP Snooping Multicast Router VLAN Configuration

IGMP Snooping Multicast Router VLAN Configuration Fields

Field	Description
<i>Interface</i>	Select the port or LAG on which to enable or disable a VLAN multicast routing interface.
<i>VLAN IDs</i>	The VLANs configured on the system that are not currently enabled as multicast router interfaces on the selected port or LAG. To enable a VLAN as a multicast router interface, click the VLAN ID to select it (or press and hold CTRL to select multiple VLAN IDs). Then, click ➡ to move the selected VLAN(s) to the <i>Configured VLAN IDs</i> field.
<i>Configured VLAN IDs</i>	The VLANs that are enabled as multicast router interfaces on the selected port or LAG. To disable a VLAN as a multicast router interface, click the VLAN ID to select it (or press and hold CTRL to select multiple VLAN IDs). Then, click ⬅ to move the selected VLAN(s) back to the <i>VLAN IDs</i> field.

Use the buttons to perform the following tasks:

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

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Configuring IGMP Snooping Querier

Use this page to configure the global IGMP snooping querier settings on the device. IGMP snooping requires that one central switch or router periodically query all end-devices on the network to announce their multicast memberships. This central device is the IGMP querier. When Layer-3 IP multicast routing protocols are enabled in a network with IP multicast routing, the IP multicast router acts as the IGMP querier. However, if the IP-multicast traffic in a VLAN needs to be Layer-2 switched only, an IP-multicast router is not required. The IGMP snooping querier can perform the IGMP snooping functions on the VLAN.

IGMP Snooping Querier Configuration

To access the *IGMP Snooping Querier Configuration* page, click **Switching > IGMP Snooping Querier > Configuration** in the navigation menu.

IGMP Snooping Querier Configuration

IGMP Snooping Querier Configuration Fields

Field	Description
<i>Admin Mode</i>	The administrative mode for the IGMP snooping querier on the device. When set to <i>Enable</i> , the IGMP snooping querier sends out periodic IGMP queries that trigger IGMP report messages from the switches that want to receive IP multicast traffic. The IGMP snooping feature listens to these IGMP reports to establish appropriate forwarding.
<i>IP Address</i>	The snooping querier address to be used as source address in periodic IGMP queries. This address is used when no IP address is configured on the VLAN on which the query is being sent.
<i>IGMP Version</i>	The IGMP protocol version used in periodic IGMP queries.
<i>Query Interval (Seconds)</i>	The amount of time the IGMP snooping querier on the device should wait between sending periodic IGMP queries.
<i>Querier Expiry Interval (Seconds)</i>	The amount of time the device remains in non-querier mode after it has discovered that there is a multicast querier in the network.

Use the buttons to perform the following tasks:

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

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

VLAN Configuration

Use this page to enable the IGMP snooping querier feature on one or more VLANs and to configure per-VLAN IGMP snooping querier settings. Only VLANs that have the IGMP snooping querier feature enabled appear in the table.

To access the *IGMP Snooping Querier VLAN Configuration* page, click **Switching** > **IGMP Snooping Querier** > **VLAN Configuration** in the navigation menu.



IGMP Snooping Querier VLAN Configuration

IGMP Snooping Querier VLAN Configuration Fields

Field	Description
VLAN ID	The VLAN on which the IGMP snooping querier is enabled. When enabling the IGMP snooping querier on a VLAN, use this menu to select the desired VLAN. Only VLANs that have been configured on the system and are not already enabled for the IGMP snooping querier appear in the menu. When modifying IGMP snooping querier settings, this field identifies the VLAN that is being configured.
Querier Election Participation	The participation mode for the IGMP snooping querier election process: Enabled The IGMP snooping querier on this VLAN participates in the querier election process when it discovers the presence of another querier in the VLAN. If the snooping querier finds that the other querier source IP address is lower than its own address, it stops sending periodic queries. If the snooping querier wins the election (because it has the lowest IP address), then it continues sending periodic queries. Disabled When the IGMP snooping querier on this VLAN sees other queriers of the same version in the VLAN, the snooping querier moves to the non-querier state and stops sending periodic queries.
Querier VLAN IP Address	The IGMP snooping querier address the VLAN uses as the source IP address in periodic IGMP queries sent on the VLAN. If this value is not configured, the VLAN uses the global IGMP snooping querier IP address.

Use the buttons to perform the following tasks:

- To enable the IGMP snooping querier feature on a VLAN, click **Add**, specify the desired settings, and click **Submit** to apply the changes.
- To change the IGMP snooping querier settings for a VLAN, select the entry to modify, click **Edit**, specify the desired settings, and click **Submit** to apply the changes.
- To disable the IGMP snooping querier feature on one or more VLANs, select each entry to change, click **Remove**, and confirm the action. Clicking this button does not remove the VLAN from the system.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

IGMP Snooping Querier VLAN Status

Use this page to view information about the IGMP snooping querier status for all VLANs that have the snooping querier enabled.

To access the *IGMP Snooping Querier VLAN Status* page, click **Switching > IGMP Snooping Querier > VLAN Status** in the navigation menu.



IGMP Snooping Querier VLAN Status

IGMP Snooping Querier VLAN Status Fields

Field	Description
<i>VLAN ID</i>	The VLAN associated with the rest of the data in the row. The table includes only VLANs that have the snooping querier enabled.
<i>State</i>	The operational state of the IGMP snooping querier on the VLAN: Querier The snooping switch is the querier in the VLAN. The snooping switch will send out periodic queries with a time interval equal to the configured querier query interval. If the snooping switch sees a better querier (numerically lower) in the VLAN, it moves to non-querier mode. Non-Querier The snooping switch is in non-querier mode in the VLAN. If the querier expiry interval timer expires, the snooping switch moves into querier mode. Disabled The snooping querier is not operational on the VLAN. The snooping querier moves to the disabled mode when IGMP snooping is not operational on the VLAN, when the querier address is not configured, or the network management address is not configured.
<i>Version</i>	The operational IGMP protocol version of the querier.
<i>Last IP Address</i>	The IP address of the last querier from which a query was snooped on the VLAN.
<i>Last Version</i>	The IGMP protocol version of the last querier from which a query was snooped on the VLAN.
<i>Max Response Time (Seconds)</i>	The maximum response time to be used in the queries that are sent by the snooping querier.

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

Creating Port Channels

Port channels, also known as link aggregation groups (LAGs), allow you to combine multiple full-duplex Ethernet links into a single logical link. Network devices treat the aggregation as if it were a single link, which increases fault tolerance and provides load sharing. You assign the port-channel (LAG) VLAN membership after you create a port-channel. The port channel by default becomes a member of the management VLAN.

A port-channel (LAG) interface can be either static or dynamic, but not both. All members of a port channel must participate in the same protocols. A static port-channel interface does not require a partner system to be able to aggregate its member ports.



Note: If you configure the maximum number of dynamic port-channels (LAGs) that your platform supports, additional port-channels that you configure are automatically static.

Static LAGs are supported. When a port is added to a LAG as a static member, it neither transmits nor receives Link Aggregation Control Protocol (LACP) Protocol Data Units (PDUs).

Port Channel Summary

Use the *Port Channel Summary* page to group one or more full duplex Ethernet links to be aggregated together to form a port-channel, which is also known as a link aggregation group (LAG). The switch can treat the port-channel as if it were a single link.

To access the page, click **Switching** > **Port Channel** > **Summary** in the navigation menu.

Name	Type	Admin Mode	STP Mode	Link State	Link Trap	Members	Admin Ports	Load Balance
et1	Static	Enabled	Enabled	Down	Enabled			Source (Destination MAC, VLAN, Ethernet, Incoming Port)
et2	Static	Enabled	Enabled	Down	Enabled			Source (Destination MAC, VLAN, Ethernet, Incoming Port)
et3	Static	Enabled	Enabled	Down	Enabled			Source (Destination MAC, VLAN, Ethernet, Incoming Port)
et4	Static	Enabled	Enabled	Down	Enabled			Source (Destination MAC, VLAN, Ethernet, Incoming Port)
et5	Static	Enabled	Enabled	Down	Enabled			Source (Destination MAC, VLAN, Ethernet, Incoming Port)
et6	Static	Enabled	Enabled	Down	Enabled			Source (Destination MAC, VLAN, Ethernet, Incoming Port)

Port Channel Summary

Port Channel Summary Fields

Field	Description
<i>Name</i>	A unique name to identify the port channel. Depending on the type of port channel, this name is automatically assigned by the system or can be configured by a system administrator.
<i>Type</i>	The type of port channel: Dynamic Uses LACP PDUs to exchange information with the link partners to help maintain the link state. To use Dynamic link aggregation on this port channel, the link partner must also support LACP. Static Does not require a partner system to be able to aggregate its member ports. When a port is added to a port channel as a static member, it neither transmits nor receives LACP PDUs. When configuring a port channel (by clicking Edit), use the <i>Static Mode</i> field to set the port channel type. If the <i>Static Mode</i> is disabled, the port channel type is Dynamic.
<i>Admin Mode</i>	The administrative mode of the port channel. When disabled, the port channel does not send and receive traffic.
<i>STP Mode</i>	The spanning tree protocol (STP) mode of the port channel. When enabled, the port channel participates in the STP operation to help prevent network loops.
<i>Link State</i>	The current link status of the port channel, which can be <i>Up</i> , <i>Up (SFP)</i> , or <i>Down</i> .
<i>Link Trap</i>	The link trap mode of the port channel. When enabled, a trap is sent to any configured SNMP receiver(s) when the link state of the port channel changes.

Port Channel Summary Fields (Continued)

Field	Description
<i>Members</i>	The ports that are members of a port channel. Each port channel can have a maximum of 8 member ports. To add ports to a port channel, select the port channel from the table and click Edit ; then, select one or more ports from the <i>Port List</i> field (press and hold CTRL to select multiple ports), click  to move the selected ports to the <i>Members</i> field, and click Submit to apply the changes.
<i>Active Ports</i>	The ports that are actively participating members of a port channel. A member port that is operationally or administratively disabled or does not have a link is not an active port.
<i>Load Balance</i>	The algorithm used to distribute traffic load among the physical ports of the port channel while preserving the per-flow packet order. The packet attributes that the load-balancing algorithm can use to determine the outgoing physical port include the following: • Source MAC, VLAN, Ethertype, Incoming Port • Destination MAC, VLAN, Ethertype, Incoming Port • Source/Destination MAC, VLAN, Ethertype, Incoming Port • Source IP and Source TCP/UDP Port Fields • Destination IP and Destination TCP/UDP Port Fields • Source/Destination IP and TCP/UDP Port Fields • Enhanced Hashing Mode

Use the buttons to perform the following tasks:

- To change the settings for an existing port channel, select the port channel from the table, click **Edit**, configure the settings as needed, and click **Submit** to apply the changes.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

Port Channel Statistics

This page displays the flap count for each port channel and their member ports. A flap occurs when a port-channel interface or port-channel member port goes down.

To access the *Port Channel Statistics* page, click **Switching > Port Channel > Statistics** in the navigation menu.

Interface	Channel Name	Type	Flap Count
3/1	ch1	Port Channel	0
3/2	ch2	Port Channel	0
3/3	ch3	Port Channel	0
3/4	ch4	Port Channel	0
3/5	ch5	Port Channel	0
3/6	ch6	Port Channel	0

Port Channel Statistics

Port Channel Statistics Fields

Field	Description
<i>Interface</i>	The port channel or member port (physical port) associated with the rest of the data in the row.
<i>Channel Name</i>	The port channel name associated with the port channel. For a physical port, this field identifies the name of the port channel of which the port is a member.
<i>Type</i>	The interface type, which is either <i>Port Channel</i> (logical link-aggregation group) or <i>Member Port</i> (physical port).
<i>Flap Count</i>	The number of times the interface has gone down. The counter for a member port is incremented when the physical port is either manually shut down by the administrator or when its link state is down. When a port channel is administratively shut down, the flap counter for the port channel is incremented, but the flap counters for its member ports are not affected. When all active member ports for a port channel are inactive (either administratively down or link down), then the port channel flap counter is incremented.

Use the buttons to perform the following tasks:

- Click **Clear Counters** to reset the flap counters for all port channels and member ports to 0.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Viewing Multicast Forwarding Database Information

The Layer-2 Multicast Forwarding Database (MFDB) is used by the switch to make forwarding decisions for packets that arrive with a multicast destination MAC address. By limiting multicasts to only certain ports in the switch, traffic is prevented from going to parts of the network where that traffic is unnecessary.

When a packet enters the switch, the destination MAC address is combined with the VLAN ID and a search is performed in the Layer-2 Multicast Forwarding Database. If no match is found, the packet is either flooded to all ports in the VLAN or discarded, depending on the switch configuration. If a match is found, the packet is forwarded only to the ports that are members of that multicast group.

Multicast Forwarding Database Summary

Use the *Multicast Forwarding Database Summary* page to view the port membership information for all active multicast address entries. The key for an entry consists of a VLAN ID and MAC address pair. Entries may contain data for more than one protocol.

To access the *Multicast Forwarding Database Summary* page, click **Switching > Multicast Forwarding Database > Summary** in the navigation menu.



Multicast Forwarding Database Summary

Multicast Forwarding Database Summary Fields

Field	Description
<i>VLAN ID</i>	The VLAN ID associated with the entry in the MFDB.
<i>MAC Address</i>	The multicast MAC address that has been added to the MFDB.
<i>Component</i>	The feature on the device that was responsible for adding the entry to the multicast forwarding database, which is one of the following: IGMP Snooping A Layer-2 feature that allows the device to dynamically add or remove ports from IPv4 multicast groups by listening to IGMP join and leave requests. MLD Snooping A Layer-2 feature that allows the device to dynamically add or remove ports from IPv6 multicast groups by listening to MLD join and leave requests. GMRP Generic Address Resolution Protocol (GARP) Multicast Registration Protocol, which helps help control the flooding of multicast traffic by keeping track of group membership information. Static Filtering A static MAC filter that was manually added to the address table by an administrator.
<i>Type</i>	The type of entry, which is one of the following: Static The entry has been manually added to the MFDB by an administrator. Dynamic The entry has been added to the MFDB as a result of a learning process or protocol.
<i>Description</i>	A text description of this multicast table entry.
<i>Interface(s)</i>	The list of interfaces that will forward or filter traffic sent to the multicast MAC address.
<i>Forwarding Interface(s)</i>	The list of forwarding interfaces. The list does not include any interfaces listed as static filtering interfaces.

To quickly find a MAC address when the list is too long to scan, enter the MAC address in the *Filter* box.

Click **Refresh** to update the information on the screen with the most current data.

Multicast Forwarding Database GMRP Table

Use the *Multicast Forwarding Database GMRP Table* page to display the entries in the multicast forwarding database (MFDB) that were added by using the GARP Multicast Registration Protocol (GMRP).

To access the page, click **Switching** > **Multicast Forwarding Database** > **GMRP** in the navigation menu.



Multicast Forwarding Database GMRP Table

Multicast Forwarding Database GMRP Table Fields

Field	Description
VLAN ID	The VLAN ID associated with the entry in the MFDB.
MAC Address	The multicast MAC address associated with the entry in the MFDB.
Type	The type of entry, which is one of the following: Static The entry has been manually added to the MFDB by an administrator. Dynamic The entry has been added to the MFDB as a result of a learning process or protocol. Entries that appear on this page have been added by using GARP.
Description	A text description of this multicast table entry.
Interface(s)	The list of interfaces that will forward or filter traffic sent to the multicast MAC address.

Click **Refresh** to update the information on the screen with the most current data.

Multicast Forwarding Database IGMP Snooping Table

This page displays the entries in the multicast forwarding database (MFDB) that were added because they were discovered by the IGMP snooping feature. IGMP snooping allows the device to dynamically add or remove ports from IPv4 multicast groups by listening to IGMP join and leave requests.

To access the page, click **Switching** > **Multicast Forwarding Database** > **IGMP Snooping** in the navigation menu.



Multicast Forwarding Database IGMP Snooping Table

Multicast Forwarding Database IGMP Snooping Table Fields

Field	Description
VLAN ID	The VLAN ID associated with the entry in the MFDB.
MAC Address	The multicast MAC address associated with the entry in the MFDB.
Type	The type of entry, which is one of the following: Static The entry has been manually added to the MFDB by an administrator. Dynamic The entry has been added to the MFDB as a result of a learning process or protocol. Entries that appear on this page have been learned by examining IGMP messages.
Description	A text description of this multicast table entry.
Interface(s)	The list of interfaces that will forward or filter traffic sent to the multicast MAC address.

Use the buttons to perform the following tasks:

- To reset the statistics to zero for all interfaces, click **Clear Counters**. You must confirm the action before the counters are reset.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

Multicast Forwarding Database Statistics

Use the *Multicast Forwarding Database Statistics* page to view statistical information about the MFDB table.

To access the *Multicast Forwarding Database Statistics* page, click **Switching** > **Multicast Forwarding Database** > **Statistics** in the navigation menu.

*Multicast Forwarding Database Statistics**Multicast Forwarding Database Statistics Fields*

Field	Description
MFDB Max Table Entries	The maximum number of entries that the multicast forwarding database can hold.
MFDB Most Entries Since Last Reset	The largest number of entries that have been present in the multicast forwarding database since the device was last reset. This value is also known as the MFDB high-water mark.
MFDB Current Entries	The current number of entries in the multicast forwarding database.

Click **Refresh** to update the information on the screen with the most current data.

Configuring Protected Ports

Use this page to configure and view protected ports groups. A port that is a member of a protected ports group is a protected port. A port that is not a member of any protected ports group is an unprotected port. Each port can be a member of only one protected ports group. Ports in the same protected ports group cannot forward traffic to other protected ports within the group, even if they are members of the same VLAN. However, a port in a protected ports group can forward traffic to ports that are in a different protected ports group. A protected port can also forward traffic to unprotected ports. Unprotected ports can forward traffic to both protected and unprotected ports.

To access the *Protected Ports Configuration* page, click **Switching > Protected Ports > Configuration** in the navigation menu.



Protected Ports Configuration

Protected Ports Configuration Fields

Field	Description
<i>Group Name</i>	This is the configured name of the protected ports group.
<i>Protected Ports</i>	The ports that are members of the protected ports group. When adding a port to a protected ports group, the <i>Available Interfaces</i> field lists the ports that are not already members of a protected ports group. To move an interface between the <i>Available Interfaces</i> and <i>Selected Interfaces</i> fields, click the port (or press and hold CTRL to select multiple ports), and then click or to move the port(s) to the desired field.

Use the buttons to perform the following tasks:

- To create a protected ports group and add ports to the group, click **Add**, configure the settings in the available fields, and click **Submit** to apply the settings.
- To change the name or the port members for an existing group, select the group to update and click **Edit**. Then, configure the settings as needed, and click **Submit** to apply the changes.
- To remove one or more protected ports groups, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Configuring Spanning Tree Protocol

The Spanning Tree Protocol (STP) provides a tree topology for any arrangement of bridges. STP also provides one path between end stations on a network, eliminating loops. Spanning tree versions supported include Common STP, Multiple STP, and Rapid STP.

Classic STP provides a single path between end stations, avoiding and eliminating loops. For information on configuring Common STP, see [“Spanning Tree CST Port Configuration” on page 164](#).

Multiple Spanning Tree Protocol (MSTP) supports multiple spanning tree instances to efficiently channel VLAN traffic over different interfaces. Each spanning tree instance behaves in the manner specified in IEEE 802.1w, Rapid Spanning Tree Protocol (RSTP), with slight modifications in the working but not the end effects (a primary effect being the rapid transitioning of the port to ‘Forwarding’). The difference between RSTP and the traditional STP (IEEE 802.1D) is the ability to configure and recognize full-duplex connectivity and ports which are connected to end stations, resulting in rapid transitioning of the port to ‘Forwarding’ state and the suppression of Topology Change Notifications (TCNs). These features are represented by the parameters ‘pointtopoint’ and ‘edgeport’. MSTP is compatible with both RSTP and STP. It behaves appropriately to STP and RSTP bridges. A MSTP bridge can be configured to behave entirely as a RSTP bridge or a STP bridge.



Note: For two bridges to be in the same region, the force version should be 802.1S and their configuration name, digest key, and revision level should match. For more information about regions and their effect on network topology, refer to the IEEE 802.1Q standard.

Spanning Tree Switch Configuration

The *Spanning Tree Switch Configuration* page contains fields for enabling STP on the switch.

To display the *Spanning Tree Switch Configuration* page, click **Switching** > **Spanning Tree** > **Switch** in the navigation menu.

Spanning Tree Switch Configuration

Spanning Tree Switch Configuration Fields

Field	Description
<i>Spanning Tree Admin Mode</i>	The STP administrative mode. If set to <i>Enable</i> , the switch participates in the root bridge election process and exchanges Bridge Protocol Data Units (BPDUs) with other switches in the spanning tree to determine the root path costs and maintain topology information.
<i>Force Protocol Version</i>	The STP version the device uses, which is one of the following: IEEE 802.1d Classic STP: Provides a single path between end stations, avoiding and eliminating loops. IEEE 802.1w Rapid Spanning Tree Protocol (RSTP): Can configure and recognize full-duplex connectivity and ports connected to end stations, for rapid transitioning of port to Forwarding state and suppression of TCNs. IEEE 802.1s Multiple Spanning Tree Protocol (MSTP): Supports multiple spanning tree instances to efficiently channel VLAN traffic over different interfaces. Compatible with RSTP and STP.

Spanning Tree Switch Configuration Fields (Continued)

Field	Description
Configuration Name	The name of the MSTP region. Each switch that participates in the same MSTP region must share the same <i>Configuration Name</i> , <i>Configuration Revision Level</i> , and MST-to-VLAN mappings.
Configuration Revision Level	The revision number of the MSTP region. This number must be the same on all switches that participate in the MSTP region.
Configuration Digest Key	A 16-byte signature of type HMAC-MD5 created from MST Configuration Table (a VLAN ID-to-MST ID mapping).
Configuration Format Selector	The version of the configuration format being used in the exchange of BPDUs.

Use the buttons to perform the following tasks:

- If you make any configuration changes, click **Submit** to apply the new settings to the switch.
- Click **Refresh** to update the information on the screen with the most current data.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

Spanning Tree CST Configuration

Use the *Spanning Tree CST Configuration* page to configure Common Spanning Tree (CST) settings. The settings and information on this page define the device within the spanning tree topology that connects all STP/RSTP bridges and MSTP regions. To access the page, click **Switching** > **Spanning Tree** > **CST** in the navigation menu.

The screenshot displays the 'Spanning Tree CST Configuration' page. At the top, there are tabs for 'Switch', 'MST', 'MST Port', 'CST' (selected), 'CST Port', and 'Statistics'. The main content area contains a list of configuration fields:

- Bridge Priority:** 0000 (0 to FFFF hex)
- Bridge Max Age:** 20 (6 to 40)
- Bridge Hello Time:** 2
- Bridge Forward Delay:** 15 (4 to 30)
- Spanning Tree Maximum Hops:** 20 (6 to 40)
- BPDU Guard:** ☐
- BPDU Filter:** ☐
- Spanning Tree Tx Hold Count:** 0 (1 to 10)
- Bridge Identifier:** 80:00:04:18:D6:31:59:F4
- Time Since Topology Change:** 1d 03:44:30
- Topology Change Count:** 0
- Topology Change:** false
- Designated Root:** 80:00:04:18:D6:31:59:F4
- Root Path Cost:** 0
- Root Port:** 00:00
- Max Age:** 20
- Forward Delay:** 15
- Hold Time:** 0
- CST Regional Root:** 80:00:04:18:D6:31:59:F4
- CST Path Cost:** 0

At the bottom of the form, there are three buttons: 'Submit', 'Refresh', and 'Cancel'. The copyright notice at the bottom right reads: '© 2013-2014 Ubiquiti Networks, Inc.'

Spanning Tree CST

Spanning Tree CST Fields

Field	Description
<i>Bridge Priority</i>	The value that helps determine which bridge in the spanning tree is elected as the root bridge during STP convergence. A lower value increases the probability that the bridge becomes the root bridge.
<i>Bridge Max Age</i>	The amount of time a bridge waits before implementing a topological change.
<i>Bridge Hello Time</i>	The amount of time the root bridge waits between sending hello BPDUs.
<i>Bridge Forward Delay</i>	The amount of time a bridge remains in a listening and learning state before forwarding packets.
<i>Spanning Tree Maximum Hops</i>	The maximum number of hops a Bridge Protocol Data Unit (BPDU) is allowed to traverse within the spanning tree region before it is discarded.
<i>BPDU Guard</i>	When enabled, <i>BPDU Guard</i> can disable edge ports that receive BPDU packets. This prevents a new device from entering the existing STP topology. Thus devices that were originally not a part of STP are not allowed to influence the STP topology.
<i>BPDU Filter</i>	When enabled, this feature filters the BPDU traffic on the edge ports. When spanning tree is disabled on a port, BPDU filtering allows BPDU packets received on that port to be dropped.
<i>Spanning Tree Tx Hold Count</i>	The maximum number of BPDUs that a bridge is allowed to send within a hello time window.
<i>Bridge Identifier</i>	A unique value that is automatically generated based on the bridge priority value and the base MAC address of the bridge. When electing the root bridge for the spanning tree, if the bridge priorities for multiple bridges are equal, the bridge with the lowest MAC address is elected as the root bridge.
<i>Time Since Topology Change</i>	The amount of time that has passed since the topology of the spanning tree has changed since the device was last reset.
<i>Topology Change Count</i>	The number of times the topology of the spanning tree has changed.
<i>Topology Change</i>	Indicates whether a topology change is in progress on any port assigned to the CST. If a change is in progress the value is <i>True</i> ; otherwise, it is <i>False</i> .
<i>Designated Root</i>	The bridge identifier of the root bridge for the CST. The identifier is made up of the bridge priority and the base MAC address.
<i>Root Path Cost</i>	The path cost to the designated root for the CST. Traffic from a connected device to the root bridge takes the least-cost path to the bridge. If the value is 0, the cost is automatically calculated based on port speed.
<i>Root Port</i>	The port on the bridge with the least-cost path to the designated root for the CST.
<i>Max Age</i>	The amount of time a bridge waits before implementing a topological change.
<i>Forward Delay</i>	The forward delay value for the root port bridge.
<i>Hold Time</i>	The minimum amount of time between transmissions of Configuration BPDUs.
<i>CST Regional Root</i>	The bridge identifier of the CST regional root. The identifier is made up of the priority value and the base MAC address of the regional root bridge.
<i>CST Path Cost</i>	The path cost to the CST tree regional root.

Use the buttons to perform the following tasks:

- If you make any configuration changes, click **Submit** to apply the new settings to the switch.
- Click **Refresh** to update the information on the screen with the most current data.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Spanning Tree CST Port Configuration

Use the CST Port page to view and configure the Common Spanning Tree (CST) settings for each interface on the device. To configure CST settings for an interface and to view additional information about the interface's role in the CST topology, select the interface to view or configure and click **Edit**.

To display the *Spanning Tree CST Port* page, click **Switching** > **Spanning Tree** > **CST Port** in the navigation menu.

Interface	Port Role	Port Forwarding State	Port Priority	Port Path Cost	Description
Q/1	Designated	Forwarding	0x0080	20000	
Q/2	Disabled	Disabled	0x0080	0	
Q/3	Disabled	Disabled	0x0080	0	
Q/4	Disabled	Disabled	0x0080	0	
Q/5	Disabled	Disabled	0x0080	0	
Q/6	Disabled	Disabled	0x0080	0	
Q/7	Disabled	Disabled	0x0080	0	
Q/8	Disabled	Disabled	0x0080	0	
Q/9	Disabled	Disabled	0x0080	0	
Q/10	Disabled	Disabled	0x0080	0	

Spanning Tree CST Port

Spanning Tree CST Port Fields

Field	Description
<i>Interface</i>	The port or link aggregation group (LAG) associated with the rest of the data in the row. When configuring CST settings for an interface, this field identifies the interface being configured.
<i>Port Role</i>	The role of the port within the CST, which is one of the following: Root A port on the non-root bridge that has the least-cost path to the root bridge. Designated A port that has the least-cost path to the root bridge on its segment. Alternate A blocked port that has an alternate path to the root bridge. Backup A blocked port that has a redundant path to the same network segment as another port on the bridge. Master The port on a bridge within an MST instance that links the MST instance to other STP regions. Disabled The port is administratively disabled and is not part of the spanning tree.
<i>Port Forwarding State</i>	Blocking The port discards user traffic and receives, but does not send, BPDUs. During the election process, all ports are in the blocking state. The port is blocked to prevent network loops. Listening The port sends and receives BPDUs and evaluates information to provide a loop-free topology. This state occurs during network convergence and is the first state in transitioning to the forwarding state. Learning The port learns the MAC addresses of frames it receives and begins to populate the MAC address table. This state occurs during network convergence and is the second state in transitioning to the forwarding state. Forwarding The port sends and receives user traffic. Disabled The port is administratively disabled and is not part of the spanning tree.

Spanning Tree CST Port Fields (Continued)

Field	Description
Port Priority	The priority for the port within the CST. This value is used in determining which port on a switch becomes the root port when two ports have the same least-cost path to the root. The port with the lower priority value becomes the root port. If the priority values are the same, the port with the lower interface index becomes the root port.
Port Path Cost	The path cost from the port to the root bridge.
Description	A user-configured description of the port. If you select an interface and click Edit , the <i>Edit CST Port Entry</i> dialog box (described below) opens and allows you to edit the CST port settings and view additional CST information for the interface.
Edit CST Port Entry dialog box – When you click Edit , this dialog box opens and allows you to configure these additional fields:	
Admin Edge Port	Select this option to administratively configure the interface as an edge port. An edge port is an interface that is directly connected to a host and is not at risk of causing a loop.
Auto-calculate Port Path Cost	Shows whether the path cost from the port to the root bridge is automatically determined by the speed of the interface (<i>Enabled</i>) or configured manually (<i>Disabled</i>).
Hello Timer	The amount of time the port waits between sending hello BPDUs.
External Port Path Cost	The cost of the path from the port to the CIST root. This value becomes important when the network includes multiple regions.
Auto-calculate External Port Path Cost	Shows whether the path cost from the port to the CIST root is automatically determined by the speed of the interface (<i>Enabled</i>) or configured manually (<i>Disabled</i>).
BPDUs Filter	Select this option to enable this feature. When enabled, this feature filters the BPDUs traffic on the edge ports. Edge ports do not need to participate in the spanning tree, so BPDUs filtering allows BPDUs packets received on edge ports to be dropped.
BPDUs Flood	Select this option to enable this feature, which determines the behavior of the interface if STP is disabled on the port and the port receives a BPDUs. If BPDUs flooding is enabled, the port will flood the received BPDUs to all the ports on the switch that are similarly disabled for spanning tree.
BPDUs Guard Effect	Shows the status (<i>Disabled</i> or <i>Enabled</i>) of BPDUs Guard Effect on the interface. When enabled, <i>BPDUs Guard Effect</i> can disable edge ports that receive BPDUs packets. This prevents a new device from entering the existing STP topology. Thus devices that were originally not a part of STP are not allowed to influence the STP topology.
Port ID	A unique value that is automatically generated based on the port priority value and the interface index.
Port Up Time Since Counters Last Cleared	The amount of time that the port has been up since the counters were cleared.
Port Mode	Used to <i>Enable</i> or <i>Disable</i> the administrative mode of spanning tree on the port.
Port Forwarding State	Blocking The port discards user traffic and receives, but does not send, BPDUs. During the election process, all ports are in the blocking state. The port is blocked to prevent network loops. Listening The port sends and receives BPDUs and evaluates information to provide a loop-free topology. This state occurs during network convergence and is the first state in transitioning to the forwarding state. Learning The port learns the MAC addresses of frames it receives and begins to populate the MAC address table. This state occurs during network convergence and is the second state in transitioning to the forwarding state. Forwarding The port sends and receives user traffic. Disabled The port is administratively disabled and is not part of the spanning tree.
Port Role	The role of the port within the CST, which is one of the following: Root A port on the non-root bridge that has the least-cost path to the root bridge. Designated A port that has the least-cost path to the root bridge on its segment. Alternate A blocked port that has an alternate path to the root bridge. Backup A blocked port that has a redundant path to the same network segment as another port on the bridge. Master The port on a bridge within an MST instance that links the MST instance to other STP regions. Disabled The port is administratively disabled and is not part of the spanning tree.
Designated Root	The bridge ID of the root bridge for the CST.
Designated Cost	The path cost offered to the LAN by the designated port.
Designated Bridge	The bridge ID of the bridge with the designated port.
Designated Port	The port ID of the designated port.

Spanning Tree CST Port Fields (Continued)

Field	Description
<i>Topology Change Acknowledge</i>	Displays <i>True</i> if the next BPDU to be transmitted for this port will have the topology change acknowledgement flag set; otherwise, displays <i>False</i> .
<i>Auto Edge</i>	When this option is selected (enabled), <i>Auto Edge</i> allows the interface to become an edge port if it does not receive any BPDUs within a given amount of time.
<i>Edge Port</i>	Displays <i>Enabled</i> if the interface is configured as an edge port; otherwise, displays <i>Disabled</i> .
<i>Point-to-point MAC</i>	Displays <i>True</i> if the link type for the interface is a point-to-point link; otherwise, displays <i>False</i> .
<i>Root Guard</i>	When this option is selected (enabled), <i>Root Guard</i> allows the interface to discard any superior information it receives to protect the root of the device from changing. The port gets put into discarding state and does not forward any frames.
<i>Loop Guard</i>	When this option is selected (enabled), <i>Loop Guard</i> prevents an interface from erroneously transitioning from blocking state to forwarding when the interface stops receiving BPDUs. The port is marked as being in loop-inconsistent state. In this state, the interface does not forward frames.
<i>TCN Guard</i>	When this option is selected (enabled), <i>TCN Guard</i> restricts the interface from propagating any topology change information received through that interface.
<i>CST Regional Root</i>	The bridge ID of the bridge that has been elected as the root bridge of the CST region.
<i>CST Path Cost</i>	The path cost from the interface to the CST regional root.
<i>Loop Inconsistent State</i>	Displays <i>True</i> if the interface is currently in a loop inconsistent state; otherwise, displays <i>False</i> . An interface transitions to a loop inconsistent state if loop guard is enabled and the port stops receiving BPDUs. In this state, the interface does not transmit frames.
<i>Transitions Into LoopInconsistent State</i>	The number of times this interface has transitioned into loop inconsistent state.
<i>Transitions Out Of LoopInconsistent State</i>	The number of times this interface has transitioned out of loop inconsistent state.

Use the buttons to perform the following tasks:

- To edit the CST port settings, click **Edit**, configure the settings as needed, and click **Submit** to apply the new settings to the switch.
- Click **Details** to display the CST port settings.
- Click **Refresh** to update the screen with most recent data.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

Spanning Tree MST Configuration

Use the *Spanning Tree MST Summary* page to view and configure the Multiple Spanning Tree Instances (MSTIs) on the device. Multiple Spanning Tree Protocol (MSTP) allows the creation of MSTIs based upon a VLAN or groups of VLANs. Configuring MSTIs creates an active topology with a better distribution of network traffic and an increase in available bandwidth when compared to classic STP.

To display the *Spanning Tree MST Port Summary* page, click **Switching** > **Spanning Tree** > **MST** in the navigation menu.

Use the buttons to perform the following tasks:

- To configure a new MSTI, click **Add** and specify the desired settings.
- To change the Priority or the VLAN associations for an existing MSTI, select the entry to modify and click **Edit**.
- To remove one or more MSTIs, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.



Spanning Tree MST Summary

Spanning Tree MST Summary Fields

Field	Description
<i>MST ID</i>	The number that identifies the MST instance.
<i>Priority</i>	The bridge priority for the spanning-tree instance. This value affects the likelihood that the bridge is selected as the root bridge. A lower value increases the probability that the bridge is selected as the root bridge.
<i># of Associated VLANs</i>	The number of VLANs that are mapped to the MSTI. This number does not contain any information about the VLAN IDs that are mapped to the instance.
<i>Bridge Identifier</i>	A unique value that is automatically generated based on the bridge priority value of the MSTI and the base MAC address of the bridge. When electing the root bridge for an MST instance, if the bridge priorities for multiple bridges are equal, the bridge with the lowest MAC address is elected as the root bridge.
<i>Time Since Topology Change</i>	The amount of time that has passed since the topology of the MSTI has changed.
<i>Designated Root</i>	The bridge identifier of the root bridge for the MST instance. The identifier is made up of the bridge priority and the base MAC address.
<i>Root Path Cost</i>	The path cost to the designated root for this MST instance. Traffic from a connected device to the root bridge takes the least-cost path to the bridge. If the value is 0, the cost is automatically calculated based on port speed.
<i>Root Port</i>	The port on the bridge with the least-cost path to the designated root for the MST instance.

Use the buttons to perform the following tasks:

- To configure a new MSTI, click **Add**, configure the settings as required, and **Submit** to apply the new settings to the switch.
- To change an existing MSTI's settings, select the entry, click **Edit**, make the required changes, and click **Submit** to apply the changes.
- To remove an MSTI, select the entry, click **Remove**, and confirm the deletion.
- Click **Refresh** to update the screen with most recent data.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Spanning Tree MST Port Configuration

Use the *Spanning Tree MST Port Summary* page to view and configure the Multiple Spanning Tree (MST) settings for each interface on the device. To configure MST settings for an interface and to view additional information about the interface's role in the MST topology, first select the appropriate MST instance from the MST ID menu. Then, select the interface to view or configure and click **Edit**.

To display the *Spanning Tree MST Port Summary* page, click **Switching > Spanning Tree > MST Port** in the navigation menu.



Note: If no MST instances have been configured on the switch, the page displays a “No MSTs Available” message and does not display the fields shown in the illustration below.



Spanning Tree MST Port Summary

Spanning Tree MST Port Summary Fields

Field	Description
<i>MST ID</i>	The menu contains the ID of each MST instance that has been created on the device.
<i>Interface</i>	The port or link aggregation group (LAG) associated with the rest of the data in the row. When configuring MST settings for an interface, this field identifies the interface being configured.
<i>Port Role</i>	The role of the port within the MST, which is one of the following: Root A port on the non-root bridge that has the least-cost path to the root bridge. Designated A port that has the least-cost path to the root bridge on its segment. Alternate A blocked port that has an alternate path to the root bridge. Backup A blocked port that has a redundant path to the same network segment as another port on the bridge. Master The port on a bridge within an MST instance that links the MST instance to other STP regions. Disabled The port is administratively disabled and is not part of the spanning tree.
<i>Port Forwarding State</i>	Blocking The port discards user traffic and receives, but does not send, BPDUs. During the election process, all ports are in the blocking state. The port is blocked to prevent network loops. Listening The port sends and receives BPDUs and evaluates information to provide a loop-free topology. This state occurs during network convergence and is the first state in transitioning to the forwarding state. Learning The port learns the MAC addresses of frames it receives and begins to populate the MAC address table. This state occurs during network convergence and is the second state in transitioning to the forwarding state. Forwarding The port sends and receives user traffic. Disabled The port is administratively disabled and is not part of the spanning tree.
<i>Port Priority</i>	The priority for the port within the MSTI. This value is used in determining which port on a switch becomes the root port when two ports have the same least-cost path to the root. The port with the lower priority value becomes the root port. If the priority values are the same, the port with the lower interface index becomes the root port.
<i>Port Path Cost</i>	The path cost from the port to the root bridge.

Spanning Tree MST Port Summary Fields (Continued)

Field	Description
<i>Description</i>	A user-configured description of the port. If you select an interface and click Edit , the <i>Edit MST Port Entry</i> dialog box (described below) opens and allows you to edit the MST port settings and view additional MST information for the interface.
<i>Edit MST Port Entry</i> dialog box – When you click Edit , this dialog box opens and allows you to configure these additional fields:	
<i>Auto-calculate Port Path Cost</i>	Shows whether the path cost from the port to the root bridge is automatically determined by the speed of the interface (Enabled) or configured manually (Disabled).
<i>Port ID</i>	A unique value that is automatically generated based on the port priority value and the interface index.
<i>Port Up Time Since Counters Last Cleared</i>	The amount of time that the port has been up since the counters were cleared.
<i>Port Mode</i>	The spanning tree administrative mode (<i>Disable</i> or <i>Enable</i>) on the port.
<i>Designated Root</i>	The bridge ID of the root bridge for the MST instance.
<i>Designated Cost</i>	The path cost offered to the LAN by the designated port.
<i>Designated Bridge</i>	The bridge ID of the bridge with the designated port.
<i>Designated Port</i>	The port ID of the designated port.
<i>Loop Inconsistent State</i>	Display <i>True</i> if the interface is currently in a loop inconsistent state; otherwise, displays <i>False</i> . An interface transitions to a loop inconsistent state if loop guard is enabled and the port stops receiving BPDUs. In this state, the interface does not transmit frames.
<i>Transitions Into LoopInconsistent State</i>	The number of times this interface has transitioned into loop inconsistent state.
<i>Transitions Out Of LoopInconsistent State</i>	The number of times this interface has transitioned out of loop inconsistent state.

Use the buttons to perform the following tasks:

- To edit the MST port settings, click **Edit**, configure the settings as needed, and click **Submit** to apply the new settings to the switch.
- Click **Details** to display the MST port settings.
- Click **Refresh** to update the screen with most recent data.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

Spanning Tree Statistics

Use the *Spanning Tree Statistics* page to view information about the number and type of bridge protocol data units (BPDUs) transmitted and received on each port.

To display the *Spanning Tree Statistics* page, click **Switching** > **Spanning Tree** > **Statistics** in the navigation menu.

Interface	STP BPDUs Rx	STP BPDUs Tx	RSTP BPDUs Rx	RSTP BPDUs Tx	MSTP BPDUs Rx	MSTP BPDUs Tx
0/1	0	0	0	0	0	40907
0/2	0	0	0	0	0	0
0/3	0	0	0	0	0	0
0/4	0	0	0	0	0	0
0/5	0	0	0	0	0	0
0/6	0	0	0	0	0	0
0/7	0	0	0	0	0	0
0/8	0	0	0	0	0	0
0/9	0	0	0	0	0	0
0/10	0	0	0	0	0	0

Spanning Tree Statistics

Spanning Tree Statistics Fields

Field	Description
<i>Interface</i>	The port or link aggregation group (LAG) associated with the rest of the data in the row.
<i>STP BPDUs Rx</i>	The number of classic STP (IEEE 802.1d) BPDUs received by the interface.
<i>STP BPDUs Tx</i>	The number of classic STP BPDUs sent by the interface.
<i>RSTP BPDUs Rx</i>	The number of RSTP (IEEE 802.1w) BPDUs received by the interface.
<i>RSTP BPDUs Tx</i>	The number of RSTP BPDUs sent by the interface.
<i>MSTP BPDUs Rx</i>	The number of MSTP (IEEE 802.1s) BPDUs received by the interface.
<i>MSTP BPDUs Tx</i>	The number of MSTP BPDUs sent by the interface.

Click **Refresh** to update the screen with most recent data.

Mapping 802.1p Priority

The IEEE 802.1p feature allows traffic prioritization at the MAC level. The switch can prioritize traffic based on the 802.1p tag attached to the L2 frame. Each port on the switch has multiple queues to give preference to certain packets over others based on the class of service (CoS) criteria you specify. When a packet is queued for transmission in a port, the rate at which it is serviced depends on how the queue is configured and possibly the amount of traffic present in the other queues of the port. If a delay is necessary, packets get held in the queue until the scheduler authorizes the queue for transmission.

Use the *802.1p Priority Mapping* page in the Class of Service submenu to assign 802.1p priority values to various traffic classes on one or more interfaces.

To display the page, click **Switching** > **Class of Service** > **802.1p Priority Mapping** in the navigation menu.

802.1p

802.1p Priority Mapping

Display **10** rows Showing 1 to 10 of 33 entries Filter:

Interface	Priority 0	Priority 1	Priority 2	Priority 3	Priority 4	Priority 5	Priority 6	Priority 7
☐ Global	1	0	0	1	2	2	3	3
☐ 0/1	1	0	0	1	2	2	3	3
☐ 0/2	1	0	0	1	2	2	3	3
☐ 0/3	1	0	0	1	2	2	3	3
☐ 0/4	1	0	0	1	2	2	3	3
☐ 0/5	1	0	0	1	2	2	3	3
☐ 0/6	1	0	0	1	2	2	3	3
☐ 0/7	1	0	0	1	2	2	3	3
☐ 0/8	1	0	0	1	2	2	3	3
☐ 0/9	1	0	0	1	2	2	3	3

First Previous 1 2 3 4 Next Last

Refresh Edit

© Copyright 2013-2014 Ubiquiti Networks, Inc.

802.1p Priority Mapping

802.1p Priority Mapping Fields

Field	Description
<i>Interface</i>	The interface associated with the rest of the data in the row. The <i>Global</i> entry represents the common settings for all interfaces, unless specifically overridden individually.
<i>Priority 0 - Priority 7</i>	The heading row lists each 802.1p priority value (<i>Priority 0</i> to <i>Priority 7</i>), and the data in the table shows which traffic class is mapped to the priority value. Incoming frames containing the designated 802.1p priority value are mapped to the corresponding traffic class in the device.
<i>Edit 802.1p Priority Mapping</i> dialog box – Click Edit to open this dialog box and configure these additional fields:	
<i>802.1p Priority</i>	The 802.1p priority value to be mapped.
<i>Traffic Class</i>	The internal traffic class to which the corresponding 802.1p priority value is mapped. The default value for each 802.1p priority level is displayed for reference.

Use the buttons to perform the following tasks:

- To edit an interface's settings, select the entry, click **Edit**, configure the settings as needed, and click **Submit** to apply the new settings to the switch.
- Click **Refresh** to update the screen with most recent data.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

Configuring Port Security

Port Security can be enabled on a per-port basis. When a port is locked, only packets with allowable source MAC addresses can be forwarded. All other packets are discarded. A MAC address can be defined as allowable by one of two methods: dynamically or statically. Note that both methods are used concurrently when a port is locked.

Dynamic locking implements a “first arrival” mechanism for Port Security. You specify how many addresses can be learned on the locked port. If the limit has not been reached, a packet with an unknown source MAC address is learned and forwarded normally. Once the limit is reached, no more addresses are learned on the port. Any packets with source MAC addresses that were not already learned are discarded. Note that you can effectively disable dynamic locking by setting the number of allowable dynamic entries to zero.

Static locking allows you to specify a list of MAC addresses that are allowed on a port. The behavior of packets is the same as for dynamic locking: only packets with an allowable source MAC address can be forwarded.

To see the MAC addresses learned on a specific port, see [“Basic Switch Configuration” on page 52](#).

Disabled ports can only be activated from the *Configuring Ports* page.

Port Security Global Administration

Use the *Port Security Global Administration* page to enable or disable the port security feature on your switch.

To access the *Port Security Global Administration* page, click **Switching > Port Security > Global** in the navigation menu.



Port Security Global Administration

Port Security Global Administration Fields

Field	Description
<i>Port Security Admin Mode</i>	The global administrative mode (<i>Enable</i> or <i>Disable</i>) for port security. The port security mode must be enabled both globally and on an interface to enforce the configured limits for the number of static and dynamic MAC addresses allowed on that interface.

Use the buttons to perform the following tasks:

- To change the *Port Security Admin Mode* setting, select **Enable** or **Disable** and click **Submit** to apply the change.
- Click **Refresh** to update the screen with most recent data.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Port Security Interface Status

Use the *Port Security Interface Status* page to configure the port security feature on a selected interface. To access the page, click **Switching** > **Port Security** > **Interface** in the navigation menu.

Interface	Port Security Mode	Max Dynamic Addresses Allowed	Max Static Addresses Allowed	Sticky Mode	Violation Trap Mode	Last Violation MAC/VLAN
G1/1	Disable	600	20	Disable	Disable	
G1/2	Disable	600	20	Disable	Disable	
G1/3	Disable	600	20	Disable	Disable	
G1/4	Disable	600	20	Disable	Disable	
G1/5	Disable	600	20	Disable	Disable	
G1/6	Disable	600	20	Disable	Disable	
G1/7	Disable	600	20	Disable	Disable	
G1/8	Disable	600	20	Disable	Disable	
G1/9	Disable	600	20	Disable	Disable	
G1/10	Disable	600	20	Disable	Disable	

Port Security Interface Status

Port Security Interface Status Fields

Field	Description
<i>Interface</i>	The interface associated with the rest of the data in the row. When configuring the port security settings for one or more interfaces, this field lists the interfaces that are being configured.
<i>Port Security Mode</i>	The administrative mode of the port security feature on the interface. The port security mode must be enabled both globally and on an interface to enforce the configured limits for the number of static and dynamic MAC addresses allowed on that interface.
<i>Max Dynamic Addresses Allowed</i>	The number of source MAC addresses that can be manually added to the port security MAC address table for an interface. If the port link goes down, the statically configured MAC addresses remain in the MAC address table. The maximum number includes all dynamically learned MAC addresses that have been converted to static MAC addresses.
<i>Max Static Addresses Allowed</i>	The number of source MAC addresses that can be manually added to the port security MAC address table for an interface. If the port link goes down, the statically configured MAC addresses remain in the MAC address table. The maximum number includes all dynamically learned MAC addresses that have been converted to static MAC addresses.
<i>Sticky Mode</i>	The sticky MAC address learning mode, which is one of the following: Enabled MAC addresses learned or manually configured on this interface are learned in sticky mode. A sticky-mode MAC address is a MAC address that does not age out and is added to the running configuration. If the running configuration is saved, the sticky addresses do not need to be relearned when the device restarts. Upon enabling sticky mode on an interface, all dynamically learned MAC addresses in the MAC address table for that interface are converted to sticky mode. Additionally, new addresses dynamically learned on the interface will also become sticky. Disabled When a link goes down on a port, all of the dynamically learned addresses are cleared from the source MAC address table for the feature. When the link is restored, the interface can once again learn addresses up to the specified limit. If sticky mode is disabled after being enabled on an interface, the sticky-mode addresses learned or manually configured on the interface are converted to dynamic entries and are automatically removed from persistent storage.
<i>Violation Trap Mode</i>	Indicates whether the port security feature sends a trap to the SNMP agent when a port is locked and a frame with a MAC address not currently in the table arrives on the port. A port is considered to be locked once it has reached the maximum number of allowed dynamic or static MAC address entries in the port security MAC address table.
<i>Last Violation MAC/VLAN</i>	The source MAC address and, if applicable, associated VLAN ID of the last frame discarded at a locked port.

Use the buttons to perform the following tasks:

- To edit the security settings of one port, select the entry, click **Edit**, configure the settings as needed, and click **Submit** to apply the changes.
- To edit the security settings for all ports, click **Edit All**, configure the settings as needed, and click **Submit** to apply the changes.
- Click **Refresh** to update the screen with most recent data.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

Port Security Statically Configured MAC Addresses

Use the *Port Security Static MAC Addresses* page to view static MAC addresses configured on an interface. From this page, you can delete statically configured MAC addresses.

To access the *Port Security Static MAC Addresses* page, click **Switching** > **Port Security** > **Static MAC** in the navigation menu.



Port Security Static MAC Addresses

Port Security Static MAC Address Fields

Field	Description
<i>Interface</i>	The interface associated with the rest of the data in the row. When adding a static MAC address entry, use the Interface menu to select the interface to associate with the permitted MAC address.
<i>Static MAC Address</i>	The MAC address of the host that is allowed to forward packets on the associated interface.
<i>VLAN ID</i>	The ID of the VLAN that includes the host with the specified MAC address.
<i>Sticky Mode</i>	Indicates whether the static MAC address entry is added in sticky mode. When adding a static MAC address entry, the <i>Sticky Mode</i> field can be selected only if it is enabled on the interface. If a static MAC address is added in sticky mode, and sticky mode is disabled on the interface, the MAC address entry is converted to a dynamic entry and will age out and be removed from the running (and saved) configuration if it is not relearned.

Use the buttons to perform the following tasks:

- To add a static MAC address, click **Add**, configure the settings as needed, and click **Submit** to apply the new settings.
- To remove a static MAC address, select the entry, click **Remove**, and confirm the deletion.
- Click **Refresh** to update the screen with most recent data.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

Port Security Dynamically Learned MAC Addresses

Use the *Port Security Dynamic MAC Addresses* page to view a table with the dynamically learned MAC addresses on an interface. With dynamic locking, MAC addresses are learned on a “first arrival” basis. You specify how many addresses can be learned on the locked port.

To access the *Port Security Dynamic MAC Addresses* page, click **Switching > Port Security > Dynamic MAC** in the navigation menu.



Port Security Dynamic MAC Addresses

Port Security Dynamic MAC Address Fields

Field	Description
<i>Interface</i>	The interface associated with the rest of the data in the row. When converting dynamic addresses to static addresses, use the Interface menu to select the interface to associate with the MAC addresses.
<i>Dynamic MAC Address</i>	The MAC address that was learned on the device. An address is dynamically learned when a frame arrives on the interface and the source MAC address in the frame is added to the MAC address table.
<i>VLAN ID</i>	The VLAN ID specified in the Ethernet frame received by the interface.

Use the buttons to perform the following tasks:

- Click **Convert to Static** to convert all MAC addresses learned on an interface to static MAC address entries. A dialog box lets you select the interface associated with the MAC address entries to convert. A static MAC address entry is written to the running configuration file and does not age out.
- Click **Refresh** to update the screen with most recent data.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Managing LLDP

The IEEE 802.1AB defined standard, Link Layer Discovery Protocol (LLDP), allows stations residing on an 802 LAN to advertise major capabilities and physical descriptions. This information is viewed by a network manager to identify system topology and detect bad configurations on the LAN.

LLDP is a one-way protocol; there are no request/response sequences. Information is advertised by stations implementing the transmit function, and is received and processed by stations implementing the receive function. The transmit and receive functions can be enabled/disabled separately per port. By default, both transmit and receive are disabled on all ports. The application is responsible for starting each transmit and receive state machine appropriately, based on the configured status and operational state of the port.

The EdgeSwitch software allows LLDP to have multiple LLDP neighbors per interface. The number of such neighbors is limited by the memory constraints. A product-specific constant defines the maximum number of neighbors supported by the switch. There is no restriction on the number of neighbors supported on a per LLDP port. If all the remote entries on the switch are filled up, the new neighbors are ignored. In case of multiple VOIP devices on a single interface, the 802.1ab component sends the Voice VLAN configuration to all the VoIP devices.

LLDP Global Configuration

Use the *LLDP Global Configuration* page to specify LLDP parameters that are applied to the switch.

To display the *LLDP Global Configuration* page, click **Switching** > **LLDP** > **Global** in the navigation menu.

LLDP Global Configuration

LLDP Global Configuration Fields

Field	Description
<i>Transmit Interval (Seconds)</i>	The number of seconds between transmissions of LLDP advertisements.
<i>Transmit Hold Multiplier (Seconds)</i>	The <i>Transmit Interval</i> multiplier value, where <i>Transmit Hold Multiplier</i> × <i>Transmit Interval</i> = the time to live (TTL) value the device advertises to neighbors.
<i>Re-Initialization Delay (Seconds)</i>	The number of seconds to wait before attempting to reinitialize LLDP on a port after the LLDP operating mode on the port changes.
<i>Notification Interval (Seconds)</i>	The minimum number of seconds to wait between transmissions of remote data change notifications to the SNMP trap receiver(s) configured on the device.

Use the buttons to perform the following tasks:

- If you make any changes to the page, click **Submit** to apply the new settings to the system.
- Click **Refresh** to update the screen with most recent data.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

LLDP Interface Configuration

Use the *LLDP Interface Summary* page to specify LLDP parameters that are applied to a specific interface.

To display the *LLDP Interface Summary* page, click **Switching > LLDP > Interface** in the navigation menu.



LLDP Interface Summary



Note: When adding or editing LLDP settings on an interface, select the appropriate check box to enable a feature, or clear the check box to disable a feature.

LLDP Interface Summary Fields

Field	Description
<i>Interface</i>	The interface associated with the rest of the data in the row. Only interfaces that have at least one LLDP setting enabled appear in the table. In the <i>Add LLDP Interface</i> window, use this field to select the interface with the LLDP settings to configure. In the <i>Edit LLDP Interface</i> window, this field identifies the interface that is being configured.
<i>Link Status</i>	The link status of the interface: <i>Up</i> or <i>Down</i> . An interface that is down does not forward traffic.
<i>Transmit</i>	The LLDP advertise (transmit) mode on the interface. If the transmit mode is enabled, the interface sends LLDP Data Units (LLDPDUs) that advertise the mandatory TLVs and any optional TLVs that are enabled.
<i>Receive</i>	The LLDP receive mode on the interface. If the receive mode is enabled, the device can receive LLDPDUs from other devices.
<i>Notify</i>	The LLDP remote data change notification status on the interface. If the notify mode is enabled, the interface sends SNMP notifications when a link partner device is added or removed.
<i>Optional TLV(s)</i>	Select each check box next to the type-length value (TLV) information to transmit. Choices include: System Name To include system name TLV in LLDP frames. To configure the System Name, see “System Description” on page 26. System Description To include system description TLV in LLDP frames. System Capabilities To include system capability TLV in LLDP frames. Port Description To include port description TLV in LLDP frames. To configure the Port Description, see “Port Description” on page 67.
<i>Transmit Management Information</i>	Select the check box to enable the transmission of management address instance. Clear the check box to disable management information transmission. The default is <i>Disabled</i> .
<i>Add/Edit LLDP Interface dialog box</i> – Click Add or Edit to open a dialog box and configure the LLDP settings for an interface:	
<i>Port Description</i>	Select this option to include the user-configured port description in the LLDPDU the interface transmits.
<i>System Name</i>	Select this option to include the user-configured system name in the LLDPDU the interface transmits. The system name, configured on the <i>System Description</i> page, is the SNMP server name for the device.
<i>System Description</i>	Select this option to include a description of the device in the LLDPDU the interface transmits. The description includes information about the product model and platform.
<i>System Capabilities</i>	Select this to advertise the primary function(s) of the device in the LLDPDU the interface transmits.

Use the buttons to perform the following tasks:

- To configure LLDP settings on an interface that does not have any LLDP settings enabled, click **Add**, configure the settings as needed, and click **Submit** to apply the new settings to the system.
- To change the LLDP settings for an interface in the table, select the entry to update, click **Edit**, configure the settings as needed, and click **Submit** to apply the new settings to the system. If you clear (disable) all LLDP settings, the entry is removed from the table.
- To clear (disable) all LLDP settings from one or more interfaces, select each entry to clear, click **Remove**, and confirm the deletion.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

LLDP Local Device Summary

Use the *LLDP Local Device Summary* page to view information about all interfaces on the device that are enabled to transmit LLDP information.

To display the page, click **Switching** > **LLDP** > **Local Devices** in the navigation menu.



LLDP Local Device Summary

LLDP Local Device Summary Fields

Field	Description
<i>Interface</i>	The interface associated with the rest of the LLDP - 802.1AB data in the row. When viewing the details for an interface, this field identifies the interface that is being viewed.
<i>Port ID</i>	The port identifier, which is the physical address associated with the interface.
<i>Port Description</i>	A description of the port. An administrator can configure this information on the <i>Port Description</i> page.
Click Details to display the following additional information about the data the interface transmits in its LLDPDUs:	
<i>Chassis ID Subtype</i>	The type of information used to identify the device in the <i>Chassis ID</i> field.
<i>Chassis ID</i>	The hardware platform identifier for the device.
<i>Port ID Subtype</i>	The type of information used to identify the interface in the <i>Port ID</i> field.
<i>System Name</i>	The user-configured system name for the device. The system name is configured on the <i>System Description</i> page and is the SNMP server name for the device.
<i>System Description</i>	The device description, which includes information about the product model and platform.
<i>System Capabilities Supported</i>	The primary function(s) the device supports.
<i>System Capabilities Enabled</i>	The primary function(s) the device supports that are enabled.
<i>Management Address</i>	The physical address associated with the management interface of the device.
<i>Management Address Type</i>	The protocol type or standard associated with the management address.

Click **Refresh** to update the information on the screen with the most current data.

Remote Device Summary

Use the *LLDP Remote Device Summary* page to view information about all interfaces on the device that are enabled to transmit LLDP information.

To display the *LLDP Remote Device Summary* page, click **Switching** > **LLDP** > **Remote Devices** in the navigation menu.



LLDP Remote Device Summary

LLDP Remote Device Summary Fields


Field	Description
<i>Interface</i>	The local interface that is enabled to receive LLDPDUs from remote devices.
<i>Remote ID</i>	The client identifier assigned to the remote system that sent the LLDPDU.
<i>Chassis ID</i>	The information the remote device sent as the Chassis ID TVL. This identifies the hardware platform for the remote system.
<i>Port ID</i>	The port on the remote system that transmitted the LLDP data.
<i>System Name</i>	The system name configured on the remote device.
Click Details to display the following additional information when the interface has received LLDPDUs from remote devices:	
Note: If the interface has not received any LLDPDUs from remote devices, a message indicates that no LLDP data has been received.	
<i>Chassis ID Subtype</i>	The type of information used to identify the device in the <i>Chassis ID</i> field.
<i>Port ID Subtype</i>	The type of information used to identify the interface in the <i>Port ID</i> field.
<i>System Description</i>	The device description, which includes information about the product model and platform.
<i>Port Description</i>	The description of the port on the remote device that transmitted the LLDP data.
<i>System Capabilities Supported</i>	The primary function(s) the remote system supports. The possible capabilities include <i>Other</i> , <i>Repeater</i> , <i>Bridge</i> , <i>WLAN AP</i> , <i>Router</i> , <i>Telephone</i> , <i>DOCSIS cable device</i> , and <i>Station</i> .
<i>System Capabilities Enabled</i>	The primary function(s) of the remote system that are both supported and enabled. The possible capabilities include <i>Other</i> , <i>Repeater</i> , <i>Bridge</i> , <i>WLAN AP</i> , <i>Router</i> , <i>Telephone</i> , <i>DOCSIS cable device</i> , and <i>Station</i> .
<i>Time To Live</i>	The number of seconds the local device should consider the LLDP data it received from the remote system to be valid.

Click **Refresh** to update the information on the screen with the most current data.

LLDP Statistics

Use the *LLDP Statistics* page to view the global and interface LLDP statistics.

To display the *LLDP Statistics* page, click **Switching > LLDP > Statistics** in the navigation menu.



LLDP Statistics

LLDP Statistics Fields

Field	Description
<i>Last Update</i>	Displays the time when an entry was created, modified, or deleted in the tables associated with the remote systems.
<i>Total Inserts</i>	Displays the number of times a complete set of information advertised by a particular MAC Service Access Point (MSAP) has been inserted into the tables associated with the remote systems.
<i>Total Deletes</i>	Displays the number of times a complete set of information advertised by a particular MAC Service Access Point (MSAP) has been deleted from the tables associated with the remote systems.
<i>Total Drops</i>	Displays the number of times a complete set of information advertised by a particular MAC Service Access Point (MSAP) could not be entered into tables associated with the remote systems because of insufficient resources.
<i>Total Ageouts</i>	Displays the number of times a complete set of information advertised by a particular MAC Service Access Point (MSAP) has been deleted from tables associated with the remote systems because the information timelines interval has expired.
<i>Interface</i>	Identifies the interfaces.
<i>Transmit Total</i>	Displays the total number of LLDP frames transmitted by the LLDP agent on the corresponding port.
<i>Receive Total</i>	Displays the total number of valid LLDP frames received by the LLDP agent on the corresponding port, while the LLDP agent is enabled.
<i>Discards</i>	Displays the number of LLDP TLVs discarded for any reason by the LLDP agent on the corresponding port.
<i>Errors</i>	Displays the number of invalid LLDP frames received by the LLDP agent on the corresponding port, while the LLDP agent is enabled.
<i>Ageouts</i>	Displays the number of age-outs that occurred on a given port. An age-out is the number of times the complete set of information advertised by a particular MAC Service Access Point (MSAP) has been deleted from tables associated with remote entries because the information timeliness interval had expired.
<i>TLV Discards</i>	Displays the number of LLDP TLVs (Type, Length, Value sets) discarded for any reason by the LLDP agent on the corresponding port.
<i>TLV Unknowns</i>	Displays the number of LLDP TLVs received on the local ports which were not recognized by the LLDP agent on the corresponding port.
<i>TLV MED</i>	Displays the total number of LLDP-MED TLVs received on the local ports.
<i>TLV 802.1</i>	Displays the total number of LLDP TLVs received on the local ports which are of type 802.1.
<i>TLV 802.3</i>	Displays the total number of LLDP TLVs received on the local ports which are of type 802.3.

Use the buttons to perform the following tasks:

- Click **Refresh** to update the page with the most current information.
- Click **Clear** to clear the LLDP statistics of all the interfaces.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

LLDP-MED

The Link Layer Discovery Protocol-Media Endpoint Discovery (LLDP-MED) is an enhancement to LLDP that features:

- Auto-discovery of LAN policies (such as VLAN, Layer-2 Priority and DiffServ settings), enabling plug and play networking.
- Device location discovery for creation of location databases.
- Extended and automated power management of Power over Ethernet endpoints.
- Inventory management, enabling network administrators to track their network devices and determine their characteristics (manufacturer, software and hardware versions, serial/asset number).

LLDP-MED Global Configuration

Use the *LLDP-MED Global Configuration* page to set global parameters for LLDP-MED operation. To display this page, click **Switching > LLDP-MED > Global** in the navigation menu.

LLDP-MED Global Configuration

LLDP-MED Global Configuration Fields

Field	Description
<i>Fast Start Repeat Count</i>	Specifies the number of LLDP-MED Protocol Data Units (PDUs) that will be transmitted when the protocol is enabled. The range is from 1 to 10. The default value is 3.
<i>Device Class</i>	Specifies local device's MED Classification. The following three represent the actual endpoints: • Class I Generic (IP Communication Controller etc.) • Class II Media (Conference Bridge etc.) • Class III Communication (IP Telephone etc.) The fourth device is Network Connectivity Device , which is typically a LAN switch or router, IEEE 802.1 bridge, or IEEE 802.11 wireless access point.

Use the buttons to perform the following tasks:

- If you make any changes on this page, click **Submit** to apply the changes.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

LLDP-MED Interface Configuration

Use the *LLDP-MED Interface Summary* page to enable LLDP-MED mode on an interface and to configure its properties. To configure the settings for one or more interfaces, select each entry to modify and click **Edit**. The same LLDP-MED settings are applied to all selected interfaces.

To display this page, click **Switching > LLDP-MED > Interface** in the navigation menu.

Interface	Link Status	MED Status	Notification Status	Operational Status	Transmit TLVs
Q/1	Up	Disable	Disable	Disable	0, 1
Q/2	Down	Disable	Disable	Disable	0, 1
Q/3	Down	Disable	Disable	Disable	0, 1
Q/4	Down	Disable	Disable	Disable	0, 1
Q/5	Down	Disable	Disable	Disable	0, 1
Q/6	Down	Disable	Disable	Disable	0, 1
Q/7	Down	Disable	Disable	Disable	0, 1
Q/8	Down	Disable	Disable	Disable	0, 1
Q/9	Down	Disable	Disable	Disable	0, 1
Q/10	Down	Disable	Disable	Disable	0, 1

LLDP-MED Interface Summary

LLDP-MED Interface Summary Fields

Field	Description
<i>Interface</i>	The interface associated with the rest of the data in the row. When configuring LLDP-MED settings, this field identifies the interfaces that are being configured.
<i>Link Status</i>	The link status of the interface, which is either <i>Up</i> or <i>Down</i> . An interface that is down does not forward traffic.
<i>MED Status</i>	The administrative status of LLDP-MED on the interface. When LLDP-MED Mode is enabled, the transmit and receive function of LLDP is effectively enabled on the interface.
<i>Notification Status</i>	Indicates whether LLDP-MED topology change notifications are enabled or disabled on the interface.
<i>Operational Status</i>	Indicates whether the interface will transmit TLVs.
<i>Transmit TLVs</i>	The LLDP-MED TLV(s) that the interface transmits: 0 Capabilities 1 Network Policy

Use the buttons to perform the following tasks:

- To add an entry to the table, click **Add**, configure the fields, and click **Submit** to apply the changes.
- To change an entry, select the entry from table, and click **Edit**. When you have completed the changes, click **Submit** to apply the changes.
- To remove an entry, select it from the table, click **Remove**, and confirm the deletion.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

LLDP-MED Local Device Information

The *LLDP-MED Local Device Summary* page displays information on LLDP-MED information advertised on the selected local interface.

To display this page, click **Switching** > **LLDP-MED** > **Local Devices** in the navigation menu.



LLDP-MED Local Device Summary

LLDP-MED Local Device Summary Fields

Field	Description
<i>Interface</i>	The interface associated with the rest of the data in the row. When viewing LLDP-MED details for an interface, this field identifies the interface that is being viewed.
<i>Port ID</i>	The MAC address of the interface. This is the MAC address that is advertised in LLDP-MED PDUs.
<i>Network Policy Information</i> – When you click Details , the <i>LLDP-MED Local Device Information</i> dialog box opens and shows the following detailed information about the LLDP-MED information the selected interface transmits.	
<i>Media Application Type</i>	The media application type transmitted in the TLV. The application types are <i>unknown</i> , <i>voicesignalling</i> , <i>guestvoice</i> , <i>guestvoicesignalling</i> , <i>softphonevoice</i> , <i>videoconferencing</i> , <i>streamingvideo</i> , and <i>videosignalling</i> . Each application type that is transmitted has the VLAN ID, priority, DSCP, tagged bit status, and unknown bit status. A port may transmit one or many such application types. This information is displayed only when a network policy TLV has been transmitted.
<i>VLAN ID</i>	The VLAN ID associated with a particular policy type.
<i>Priority</i>	The user priority associated with a particular policy type.
<i>DSCP</i>	The DSCP value associated with a particular policy type.
<i>Unknown Bit Status</i>	The unknown bit associated with a particular policy type.
<i>Tagged Bit Status</i>	Identifies whether the network policy is defined for tagged or untagged VLANs.
<i>Location Information:</i>	
<i>Sub Type</i>	The type of location information: Coordinate Based The location map coordinates (latitude, longitude and altitude) of the device. Civic Address The civic or street address location of the device. ELIN The Emergency Call Service (ECS) Emergency Location Identification Number (ELIN) of the device.
<i>Information</i>	This column displays the information related to the coordinates, civic address, and ELIN for the device.

Use the buttons to perform the following tasks:

- To display detailed information about an LLDP-MED interface, select the interface and click **Details**. A window displays the fields in the *Network Policy Information* section of the table below.
- Click **Refresh** to refresh the page with the most current data from the switch.

LLDP-MED Remote Device Information

The *LLDP-MED Remote Device Summary* page displays information about the remote devices the local system has learned about through the LLDP-MED data units received on its interfaces. Information is available about remote devices only if an interface receives an LLDP-MED data unit from a device.

To display this page, click **Switching > LLDP-MED > Remote Devices** in the navigation menu.



LLDP-MED Remote Device Summary

LLDP-MED Remote Device Summary Fields

Field	Description
<i>Interface</i>	The local interface that has received LLDP-MED data units from remote devices.
<i>Remote ID</i>	The client identifier assigned to the remote system that sent the LLDP-MED data unit.
<i>Capability Information:</i>	
<i>Supported Capabilities</i>	The supported capabilities that were received in the MED TLV on this interface.
<i>Enabled Capabilities</i>	The supported capabilities on the remote device that are also enabled.
<i>Device Class</i>	The MED Classification advertised by the TLV from the remote device. The following three classifications represent the actual endpoints: Class I Generic (for example, IP Communication Controller) Class II Media (for example, Conference Bridge) Class III Communication (for example, IP Telephone) The fourth device is Network Connectivity Device, which is typically a device such as a LAN switch or router, IEEE 802.1 bridge, or IEEE 802.11 wireless access point.
<i>Network Policy Information:</i>	
<i>Media Application Type</i>	The media application type received in the TLV from the remote device. The application types are <i>unknown</i> , <i>voicesignalling</i> , <i>guestvoice</i> , <i>guestvoicesignalling</i> , <i>softphonevoice</i> , <i>videoconferencing</i> , <i>streamingvideo</i> , and <i>videosignalling</i> . Each application type that is transmitted has the VLAN ID, priority, DSCP, tagged bit status, and unknown bit status. The port on the remote device may transmit one or many such application types. This information is displayed only when a network policy TLV has been received.
<i>VLAN ID</i>	The VLAN ID associated with a particular policy type.
<i>Priority</i>	The user priority associated with a particular policy type.
<i>DSCP</i>	The DSCP value associated with a particular policy type.
<i>Unknown Bit Status</i>	The unknown bit associated with a particular policy type.
<i>Tagged Bit Status</i>	Identifies whether the network policy is defined for tagged or untagged VLANs.
<i>Inventory Information:</i>	
<i>Hardware Revision</i>	The hardware version advertised by the remote device.
<i>Firmware Revision</i>	The firmware version advertised by the remote device.
<i>Software Revision</i>	The software version advertised by the remote device.
<i>Serial Number</i>	The serial number advertised by the remote device.

LLDP Remote Device Information Fields (Continued)

Field	Description
<i>Manufacturer Name</i>	The name of the system manufacturer advertised by the remote device.
<i>Model Name</i>	The name of the system model advertised by the remote device.
<i>Asset ID</i>	The system asset ID advertised by the remote device.
<i>Location Information:</i>	
<i>Sub Type</i>	The type of location information advertised by the remote device.
<i>Information</i>	The text description of the location information included in the subtype.
<i>Extended PoE</i>	Indicates whether the remote device is advertised as a PoE device.
<i>Device Type</i>	If the remote device is a PoE device, this field identifies the PoE device type of the remote device connected to this port.

Use the buttons to perform the following tasks:

- To view additional information about a remote device, select the interface that received the LLDP-MED data and click **Details**. The *LLDP-MED Remote Device Information* window appears and displays the fields in the table below.
- Click **Refresh** to refresh the page with the most current data from the switch.

Chapter 5: Configuring Routing

EdgeSwitch supports IP routing. Use the links in the Routing navigation menu folder to manage routing on the system. This section contains the following information:

- **[“Configuring ARP” on page 188](#)**
- **[“Configuring Global IP Settings” on page 191](#)**
- **[“Router” on page 200](#)**
- **[“Configuring Policy-Based Routing” on page 203](#)**

When a packet enters the switch, the destination MAC address is checked to see if it matches any of the configured routing interfaces. If it does, then the silicon searches the host table for a matching destination IP address. If an entry is found, then the packet is routed to the host. If there is not a matching entry, then the switch performs a longest prefix match on the destination IP address. If an entry is found, then the packet is routed to the next hop. If there is no match, then the packet is routed to the next hop specified in the default route. If there is no default route configured, then the packet is passed to the software to be handled appropriately.

The routing table can have entries added either statically by the administrator or dynamically via a routing protocol. The host table can have entries added either statically by the administrator or dynamically via ARP.

Configuring ARP

The Address Resolution Protocol (ARP) associates a Layer-2 MAC address with a Layer-3 IPv4 address. The EdgeSwitch software features both dynamic and manual ARP configuration. With manual ARP configuration, you can statically add entries into the ARP table.

ARP is a necessary part of the internet protocol (IP) and is used to translate an IP address to a media (MAC) address, defined by a local area network (LAN) such as Ethernet. A station needing to send an IP packet must learn the MAC address of the IP destination, or of the next hop router, if the destination is not on the same subnet. This is achieved by broadcasting an ARP request packet, to which the intended recipient responds by unicasting an ARP reply containing its MAC address. Once learned, the MAC address is used in the destination address field of the Layer-2 header prepended to the IP packet.

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

The number of supported ARP entries is platform-dependent.

Devices can be moved in a network, which means the IP address that was at one time associated with a certain MAC address is now found using a different MAC, or may have disappeared from the network altogether (i.e., it has been reconfigured, disconnected, or powered off). This leads to stale information in the ARP cache unless entries are updated in reaction to new information seen on the network, periodically refreshed to determine if an address still exists, or removed from the cache if the entry has not been identified as a sender of an ARP packet during the course of an ageout interval, usually specified via configuration.

The **Routing > ARP Table** submenu contains links to the following UI pages that configure and display ARP-related details:

- **[“ARP Table” on page 189](#)**
- **[“ARP Table Configuration” on page 190](#)**

ARP Table

Use the *ARP Table* page to add an entry to the Address Resolution Protocol table.

To display the page, click **Routing > ARP Table > Summary** in the navigation menu. The *ARP Table* is displayed at the bottom of the page, and contains the fields in the table below.



ARP Table

ARP Table Fields

Field	Description
<i>IP Address</i>	The IP address of a network host on a subnet attached to one of the device's routing interfaces. When adding a static ARP entry, specify the IP address for the entry after you click Add .
<i>MAC Address</i>	The unicast MAC address (hardware address) associated with the network host. When adding a static ARP entry, specify the MAC address to associate with the IP address in the entry.
<i>Interface</i>	The routing interface associated with the ARP entry. The network host is associated with the device through this interface.
<i>Type</i>	The ARP entry type: Dynamic An ARP entry that has been learned by the router Gateway A dynamic ARP entry that has the IP address of a routing interface Local An ARP entry associated with the MAC address of a routing interface on the device Static An ARP entry configured by the user
<i>Age</i>	The age of the entry since it was last learned or refreshed. This value is specified for <i>Dynamic</i> or <i>Gateway</i> entry types only (it is left blank for all other entry types).

Use the buttons to perform the following tasks:

- To add a static ARP entry, click **Add**. In the *Add Static ARP Entry* dialog box, configure the information for the new ARP entry, and click **Submit** to apply the changes and save the entry to the ARP table.
- To delete one or more ARP entries, select each entry to delete and click **Remove**. Note that ARP entries designated as *Local* cannot be removed.
- Click **Refresh** to update the information on the screen with the most current data.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

ARP Table Configuration

Use the *ARP Table Configuration* page to change the configuration parameters for the ARP table. You can also use this screen to display the contents of the table.

To display the page, click **Routing > ARP Table > Configuration** in the navigation menu.

ARP Table Configuration

ARP Table Configuration Fields

Field	Description
<i>Age Time</i>	The amount of time, in seconds, that a dynamic ARP entry remains in the ARP table before aging out.
<i>Response Time</i>	The amount of time, in seconds, that the device waits for an ARP response to an ARP request that it sends.
<i>Retries</i>	The maximum number of times an ARP request will be retried after an ARP response is not received. The number includes the initial ARP request.
<i>Cache Size</i>	The maximum number of entries allowed in the ARP table. This number includes all static and dynamic ARP entries.
<i>Dynamic Renew</i>	When selected, this option allows the ARP component to automatically attempt to renew dynamic ARP entries when they age out.

Use the buttons to perform the following tasks:

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

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Configuring Global IP Settings

The **Routing > IP** folder contains links to the following UI pages that configure and display IP routing data:

- [“Routing IP Configuration” on page 191](#)
- [“Routing IP Interface Summary” on page 193](#)
- [“Routing IP Interface Configuration” on page 195](#)
- [“Routing IP Statistics” on page 197](#)

Routing IP Configuration

Use the *Routing IP Configuration* page to configure global routing settings on the device. Routing provides a means of transmitting IP packets between subnets on the network. Routing configuration is necessary only if the device is used as a Layer-3 device that routes packets between subnets. If the device is used as a Layer-2 device that handles switching only, it typically connects to an external Layer-3 device that handles the routing functions; therefore, routing configuration is not required on the Layer-2 device.

To display the page, click **Routing > IP > Configuration** in the navigation menu.

Routing IP Configuration

Routing IP Configuration Fields

Field	Description
<i>Routing Mode</i>	The administrative mode of routing on the device. The options are as follows: • Enable The device can act as a Layer-3 device by routing packets between interfaces configured for IP routing. • Disable The device acts as a Layer-2 bridge and switches traffic between interfaces. The device does not perform any internetwork routing.
<i>ICMP Echo Replies</i>	Select <i>Enable</i> or <i>Disable</i> from the drop-down menu. If you select <i>Enable</i> , then only the router can send ECHO replies. By default, ICMP Echo Replies are sent for echo requests.
<i>ICMP Redirects</i>	Select this option to allow the device to send ICMP Redirect messages to hosts. An ICMP Redirect message notifies a host when a better route to a particular destination is available on the network segment.
<i>ICMP Rate Limit Interval</i>	To control the ICMP error packets, you can specify the number of ICMP error packets that are allowed per burst interval. By default, the rate limit is 100 packets per second, i.e. the burst interval is 1000 milliseconds. To disable ICMP rate limiting, set this field to zero. The valid rate interval range is 0 to 2147483647 milliseconds.

Routing IP Configuration Fields (Continued)

Field	Description
<i>ICMP Rate Limit Burst Size</i>	To control the ICMP error packets, you can specify the number of ICMP error packets that are allowed per burst interval. By default, the burst size is 100 packets. When the burst interval is zero, then configuring this field is not a valid option. The valid burst size range is 1 to 200.
<i>Static Route Preference</i>	The default distance (preference) for static routes. Lower route-distance values are preferred when determining the best route. The value configured for <i>Static Route Preference</i> is used when using the CLI to configure a static route and no preference is specified. Changing the <i>Static Route Preference</i> does not update the preference of existing static routes.
<i>Local Route Preference</i>	The default distance (preference) for local routes.
<i>Maximum Next Hops</i>	The maximum number of hops supported by the switch. This is a read-only value.
<i>Maximum Routes</i>	The maximum number of routes (routing table size) supported by the switch.
<i>Global Default Gateway</i>	The IP address of the default gateway for the device. If the destination IP address in a packet does not match any routes in the routing table, the packet is sent to the default gateway. The gateway specified in this field is more preferred than a default gateway learned from a DHCP server. Use the buttons next to this field as follows:  Click this button to configure the default gateway.  Click this button to reset the IP address of the default gateway to the factory default value.

Use the buttons to perform the following tasks:

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

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Routing IP Interface Summary

The *Routing IP Interface Summary* page shows summary information about the routing configuration for all interfaces. To view additional routing configuration information for an interface, select the interface with the settings to view and click **Details**.

To display the page, click **Routing > IP > Interface Summary** in the navigation menu.

Interface	Status	IP Address	Subnet Mask	Admin Mode	State	MAC Address	Proxy ARP	IP MTU
0/1	Down	0.0.0.0	0.0.0.0	Enabled	Active	04:18:D6:31:59:F5	Disabled	1500
0/2	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	04:18:D6:31:59:F5	Disabled	1500
0/3	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	04:18:D6:31:59:F5	Disabled	1500
0/4	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	04:18:D6:31:59:F5	Disabled	1500
0/5	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	04:18:D6:31:59:F5	Disabled	1500
0/6	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	04:18:D6:31:59:F5	Disabled	1500
0/7	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	04:18:D6:31:59:F5	Disabled	1500
0/8	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	04:18:D6:31:59:F5	Disabled	1500
0/9	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	04:18:D6:31:59:F5	Disabled	1500
0/10	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	04:18:D6:31:59:F5	Disabled	1500

Routing IP Interface Summary

Routing IP Interface Summary Fields

Field	Description
<i>Interface</i>	The interface associated with the rest of the data in the row. When viewing details about the routing settings for an interface, this field identifies the interface being viewed.
<i>Status</i>	Indicates whether the interface is capable of routing IP packets (<i>Up</i>) or cannot route packets (<i>Down</i>). For the status to be Up, the routing mode and administrative mode for the interface must be enabled. Additionally, the interface must have an IP address and be physically up (active link).
<i>IP Address</i>	The IP address of the interface.
<i>Subnet Mask</i>	The IP subnet mask for the interface (also known as the network mask or netmask). It defines the portion of the interface's IP address that is used to identify the attached network.
<i>Admin Mode</i>	The administrative mode of the interface, which is either <i>Enabled</i> or <i>Disabled</i> .
<i>State</i>	The state of the interface, which is either <i>Active</i> or <i>Inactive</i> . An interface is considered active if the link is up, and the interface is in a forwarding state.
<i>MAC Address</i>	The burned-in physical address of the interface. The format is six two-digit hexadecimal numbers separated by colons, for example 00:06:29:32:81:40.
<i>Proxy ARP</i>	Indicates whether proxy ARP is enabled or disabled on the interface. When proxy ARP is enabled, the interface can respond to an ARP request for a host other than itself. An interface can act as an ARP proxy if it is aware of the destination and can route packets to the intended host, which is on a different subnet than the host that sent the ARP request.
<i>IP MTU</i>	The largest IP packet size the interface can transmit, in bytes. The IP Maximum Transmission Unit (MTU) is the maximum frame size minus the length of the Layer-2 header.

Routing IP Interface Summary Fields (Continued)

Field	Description
<i>Details</i> window – If you select an interface and click Details , the <i>Details</i> window opens and displays the following additional routing information for the selected interface:	
<i>Routing Mode</i>	Indicates whether routing is administratively <i>Enabled</i> or <i>Disabled</i> on the interface.
<i>Link Speed Data Rate</i>	The physical link data rate of the interface.
<i>IP Address Configuration Method</i>	The source of the IP address, which is one of the following: • None The interface does not have an IP address. • Manual The IP address has been statically configured by an administrator. • DHCP The IP address has been learned dynamically through DHCP. If the method is DHCP but the interface does not have an IP address, the interface is unable to acquire an address from a network DHCP server.
<i>Bandwidth</i>	The configured bandwidth on this interface. This setting communicates the speed of the interface to higher-level protocols.
<i>Encapsulation Type</i>	The link layer encapsulation type for packets transmitted from the interface, which can be either <i>Ethernet</i> or <i>SNAP</i> .
<i>Forward Net Directed Broadcasts</i>	Indicates how the interface handles network-directed broadcast packets. A network-directed broadcast is a broadcast directed to a specific subnet. The possible values are as follows: • Enabled Network directed broadcasts are forwarded. • Disabled Network directed broadcasts are dropped.
<i>Local Proxy ARP</i>	Indicates whether local proxy ARP is <i>Enabled</i> or <i>Disabled</i> on the interface. When local proxy ARP is enabled, the interface can respond to an ARP request for a host other than itself. Unlike proxy ARP, local proxy ARP allows the interface to respond to ARP requests for a host that is on the same subnet as the host that sent the ARP request. This feature is useful when a host is not permitted to reply to an ARP request from another host in the same subnet, for example when using the protected ports feature.
<i>Destination Unreachables</i>	Displays <i>Enabled</i> if the interface is allowed to send ICMP Destination Unreachable message to a host if the intended destination cannot be reached for some reason. If the field displays <i>Disabled</i> , this interface will not send ICMP Destination Unreachable messages to inform the host about the error in reaching the intended destination.
<i>ICMP Redirects</i>	Displays <i>Enabled</i> if the interface is allowed to send ICMP Redirect messages; otherwise, displays <i>Disabled</i> . The device sends an ICMP Redirect message on an interface only if ICMP Redirects are enabled both globally and on the interface. An ICMP Redirect message notifies a host when a better route to a particular destination is available on the network segment.

Use the buttons to perform the following tasks:

- To edit an interface's routing configuration, select the interface and click **Edit**. The display changes to the *Routing IP Interface Configuration* page; for instructions on using this page, see [“Routing IP Interface Configuration” on page 195](#).
- To view detailed routing information on an interface, select the interface's entry and click **Details**.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Routing IP Interface Configuration

Use the *Routing IP Interface Configuration* page to configure the IP routing settings for each interface.

To display the page, click **Routing > IP > Interface Configuration** in the navigation menu.

Routing IP Interface Configuration

Routing IP Interface Configuration Fields

Field	Description
<i>Interface</i>	The menu contains all interfaces that can be configured for routing. To configure routing settings for an interface, select it from the menu and then configure the rest of the settings on the page.
<i>Status</i>	Indicates whether the interface is currently capable of routing IP packets (<i>Up</i>) or cannot route packets (<i>Down</i>). For the status to be <i>Up</i> , the routing mode and administrative mode for the interface must be enabled. Additionally, the interface must have an IP address and be physically up (active link).
<i>Routing Mode</i>	Used to <i>Enable</i> or <i>Disable</i> the administrative mode of IP routing on the interface.

Routing IP Interface Configuration Fields (Continued)

Field	Description
<i>Admin Mode</i>	The administrative mode of the interface. If set to <i>Disable</i> , the interface cannot forward traffic.
<i>State</i>	The state of the interface, which is either <i>Active</i> or <i>Inactive</i> . An interface is considered active if the link is up, and the interface is in a forwarding state.
<i>Link Speed Data Rate</i>	The physical link data rate of the interface.
<i>IP Address Configuration Method</i>	The method to use for configuring an IP address on the interface, which can be one of the following: None No address is to be configured. Manual The address is to be statically configured. When this option is selected you can specify the IP address and subnet mask in the available fields. DHCP The interface will attempt to acquire an IP address from a network DHCP server.
<i>IP Address</i>	The IP address of the interface. This field can be configured only when the selected <i>IP Address Configuration Method</i> is <i>Manual</i> . If the method is <i>DHCP</i> , the interface attempts to lease an IP address from a DHCP server on the network, and the IP address appears in this field (read-only) after it is acquired. If this field is blank, the <i>IP Address Configuration Method</i> might be <i>None</i> , or the method might be <i>DHCP</i> and the interface is unable to lease an address.
<i>Subnet Mask</i>	The IP subnet mask for the interface (also known as the network mask or netmask). This field can be configured only when the selected <i>IP Address Configuration Method</i> is <i>Manual</i> .
<i>MAC Address</i>	The burned-in physical address of the interface. The format is six two-digit hexadecimal numbers separated by colons, for example 00:06:29:32:81:40.
<i>IP MTU</i>	The largest IP packet size the interface can transmit, in bytes. The IP Maximum Transmission Unit (MTU) is the maximum frame size minus the length of the Layer-2 header. Click  to reset this field.
<i>Bandwidth</i>	The configured bandwidth on this interface. This setting communicates the speed of the interface to higher-level protocols.
<i>Encapsulation Type</i>	The link layer encapsulation type for packets transmitted from the interface: <i>Ethernet</i> or <i>SNAP</i> .
<i>Forward Net Directed Broadcasts</i>	Determines how the interface handles network-directed broadcast packets. A network-directed broadcast is a broadcast directed to a specific subnet. If this option is selected, network directed broadcasts are forwarded. If this option is cleared, network directed broadcasts are dropped.
<i>Proxy ARP</i>	When this option is selected, proxy ARP is enabled, and the interface can respond to an ARP request for a host other than itself. An interface can act as an ARP proxy if it is aware of the destination and can route packets to the intended host, which is on a different subnet than the host that sent the ARP request.
<i>Local Proxy ARP</i>	When this option is selected, local proxy ARP is enabled, and the interface can respond to an ARP request for a host other than itself. Unlike proxy ARP, local proxy ARP allows the interface to respond to ARP requests for a host that is on the same subnet as the host that sent the ARP request. This feature is useful when a host is not permitted to reply to an ARP request from another host in the same subnet, for example when using the protected ports feature.
<i>Destination Unreachables</i>	When this option is selected, the interface is allowed to send ICMP Destination Unreachable message to a host if the intended destination cannot be reached for some reason. If this option is clear, the interface will not send ICMP Destination Unreachable messages to inform the host about the error in reaching the intended destination.
<i>ICMP Redirects</i>	When this option is selected, the interface is allowed to send ICMP Redirect messages. The device sends an ICMP Redirect message on an interface only if ICMP Redirects are enabled both globally and on the interface. An ICMP Redirect message notifies a host when a better route to a particular destination is available on the network segment.
<i>Secondary IP Address</i>	To add a secondary IP address on the interface, click  in the header row and enter the address in the appropriate field in the <i>Secondary IP Address Configuration</i> window. You can add one or more secondary IP addresses to an interface only if the interface already has a primary IP address. To remove a configured secondary IP address, click the entry's associated  button. To remove <i>all</i> configured secondary IP addresses, click  in the header row.
<i>Secondary Subnet Mask</i>	The subnet mask associated with the secondary IP address. You configure this field in the <i>Secondary IP Address Configuration</i> window.

Use the buttons to perform the following tasks:

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

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Routing IP Statistics

The statistics reported on the *Routing IP Statistics* page are as specified in RFC 1213.

To display the page, click **Routing > IP > Statistics** in the navigation menu.

Configuration	Interface Summary	Interface Configuration	Loopback Configuration	Statistics
Routing IP Statistics				
ipInReceives				138884
ipInHdrErrors				0
ipInAddrErrors				0
ipFwdDiscards				0
ipInUnknownProtos				0
ipInDiscards				0
ipInDelivers				138884
ipOutRequests				30882
ipOutDiscards				0
ipOutNoRoutes				0
ipReasmTimeout				0
ipReasmFails				0
ipReasmOKs				0
ipReasmFails				0
ipFragOKs				0
ipFragFails				0
ipFragCreates				0
ipRoutingDiscards				0
icmpInMgts				4
icmpInErrors				0
icmpInDestUnreachs				1
icmpInTimeExcds				0
icmpInParamProbs				0
icmpInSrcQuenchs				0
icmpInBadlenets				0
icmpInEchoes				1
icmpInEchoReps				0
icmpInTimeMgmts				0
icmpInTimeMgmtReps				0
icmpInAddMasks				0
icmpInAddMaskReps				0
icmpOutMgts				4
icmpOutErrors				0
icmpOutDestUnreachs				1
icmpOutTimeExcds				0
icmpOutParamProbs				0
icmpOutSrcQuenchs				0
icmpOutBadlenets				0
icmpOutEchoes				0
icmpOutEchoReps				1
icmpOutTimeMgmts				0
icmpOutTimeMgmtReps				0
icmpOutAddMasks				0
Refresh				

Routing IP Statistics

Routing IP Statistics Fields

Field	Description
<i>IpInReceives</i>	The total number of input datagrams received from interfaces, including those received in error.
<i>IpInHdrErrors</i>	The number of input datagrams discarded due to errors in their IP headers, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, errors discovered in processing their IP options, etc.
<i>IpInAddrErrors</i>	The number of input datagrams discarded because the IP address in their IP header's destination field was not a valid address to be received at this entity. This count includes invalid addresses (e.g., 0.0.0.0) and addresses of unsupported Classes (e.g., Class E). For entities which are not IP Gateways and therefore do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address.
<i>IpForwDatagrams</i>	The number of input datagrams for which this entity was not their final IP destination, as a result of which an attempt was made to find a route to forward them to that final destination. In entities which do not act as IP Gateways, this counter includes only those packets which were Source-Routed via this entity, and the Source-Route option processing was successful.
<i>IpInUnknownProtos</i>	The number of locally-addressed datagrams received successfully but discarded because of an unknown or unsupported protocol.
<i>IpInDiscards</i>	The number of input IP datagrams for which no problems were encountered to prevent their continued processing, but which were discarded (e.g., for lack of buffer space). Note that this counter does not include any datagrams discarded while awaiting re-assembly.
<i>IpInDelivers</i>	The total number of input datagrams successfully delivered to IP user-protocols (including ICMP).
<i>IpOutRequests</i>	The total number of IP datagrams which local IP user-protocols (including ICMP) supplied to IP in requests for transmission. Note that this counter does not include any datagrams counted in <i>IpForwDatagrams</i> .
<i>IpOutDiscards</i>	The number of output IP datagrams for which no problem was encountered to prevent their transmission to their destination, but which were discarded (e.g., for lack of buffer space). Note that this counter would include datagrams counted in <i>IpForwDatagrams</i> if any such packets met this (discretionary) discard criterion.
<i>IpOutNoRoutes</i>	The number of IP datagrams discarded because no route could be found to transmit them to their destination. Note that this counter includes any packets counted in <i>IpForwDatagrams</i> which meet this 'no-route' criterion. Note that this includes any datagrams which a host cannot route because all of its default gateways are down.
<i>IpReasmTimeout</i>	The maximum number of seconds which received fragments are held while they are awaiting reassembly at this entity.
<i>IpReasmReqds</i>	The number of IP fragments received which needed to be reassembled at this entity.
<i>IpReasmOKs</i>	The number of IP datagrams successfully re-assembled.
<i>IpReasmFails</i>	The number of failures detected by the IP re-assembly algorithm (for whatever reason: timed out, errors, etc.). Note that this is not necessarily a count of discarded IP fragments since some algorithms can lose track of the number of fragments by combining them as they are received.
<i>IpFragOKs</i>	The number of IP datagrams that have been successfully fragmented at this entity.
<i>IpFragFails</i>	The number of IP datagrams that have been discarded because they needed to be fragmented at this entity but could not be, e.g., because their Don't Fragment flag was set.
<i>IpFragCreates</i>	The number of IP datagram fragments that have been generated as a result of fragmentation at this entity.
<i>IpRoutingDiscards</i>	The number of routing entries which were chosen to be discarded even though they are valid. One possible reason for discarding such an entry could be to free-up buffer space for other routing entries.
<i>IcmpInMsgs</i>	The total number of ICMP messages which the entity received. Note that this counter includes all those counted by <i>IcmpInErrors</i> .
<i>IcmpInErrors</i>	The number of ICMP messages which the entity received but determined as having ICMP-specific errors (bad ICMP checksums, bad length, etc.).
<i>IcmpInDestUnreachs</i>	The number of ICMP Destination Unreachable messages received.
<i>IcmpInTimeExcds</i>	The number of ICMP Time Exceeded messages received.
<i>IcmpInParmProbs</i>	The number of ICMP Parameter Problem messages received.
<i>IcmpInSrcQuenchs</i>	The number of ICMP Source Quench messages received.
<i>IcmpInRedirects</i>	The number of ICMP Redirect messages received.
<i>IcmpInEchos</i>	The number of ICMP Echo (request) messages received.

Routing IP Statistics Fields (Continued)

Field	Description
<i>IcmpInEchoReps</i>	The number of ICMP Echo Reply messages received.
<i>IcmpInTimestamps</i>	The number of ICMP Timestamp (request) messages received.
<i>IcmpInTimestampReps</i>	The number of ICMP Timestamp Reply messages received.
<i>IcmpInAddrMasks</i>	The number of ICMP Address Mask Request messages received.
<i>IcmpInAddrMaskReps</i>	The number of ICMP Address Mask Reply messages received.
<i>IcmpOutMsgs</i>	The total number of ICMP messages which this entity attempted to send. Note that this counter includes all those counted by <i>IcmpOutErrors</i> .
<i>IcmpOutErrors</i>	The number of ICMP messages which this entity did not send due to problems discovered within ICMP such as a lack of buffers. This value should not include errors discovered outside the ICMP layer such as the inability of IP to route the resultant datagram. In some implementations there may be no types of error which contribute to this counter's value.
<i>IcmpOutDestUnreachs</i>	The number of ICMP Destination Unreachable messages sent.
<i>IcmpOutTimeExcds</i>	The number of ICMP Time Exceeded messages sent.
<i>IcmpOutParmProbs</i>	The number of ICMP Parameter Problem messages sent.
<i>IcmpOutSrcQuenchs</i>	The number of ICMP Source Quench messages sent.
<i>IcmpOutRedirects</i>	The number of ICMP Redirect messages sent. For a host, this object is always zero, since hosts do not send redirects.
<i>IcmpOutEchos</i>	The number of ICMP Echo (request) messages sent.
<i>IcmpOutEchoReps</i>	The number of ICMP Echo Reply messages sent.
<i>IcmpOutTimestamps</i>	The number of ICMP Timestamp (request) messages.
<i>IcmpOutTimestampReps</i>	The number of ICMP Timestamp Reply messages sent.
<i>IcmpOutAddrMasks</i>	The number of ICMP Address Mask Request messages sent.

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

Router

The **Routing** > **Router** menu contains links to UI pages that configure and display route tables.

Route Table

The route table manager collects routes from multiple sources: static routes and local routes. The route table manager may learn multiple routes to the same destination from multiple sources. The route table lists all routes. The best routes table displays only the most preferred route to each destination.

To display the *Route Table Summary* page, click **Routing** > **Router** > **Route Table** in the navigation menu.



Route Table Summary

Route Table Summary Fields


Field	Description
<i>Network Address</i>	The IP route prefix for the destination.
<i>Subnet Mask</i>	Also referred to as the subnet/network mask, this indicates the portion of the IP interface address that identifies the attached network.
<i>Protocol</i>	This field tells which protocol created the specified route. A route can be created in the following ways: • Dynamically learned through a supported routing protocol • Dynamically learned by being a directly-attached local route • Statically configured by an administrator • Configured as a default route by an administrator
<i>Next Hop IP Address</i>	The outgoing router IP address to use when forwarding traffic to the next router (if any) in the path towards the destination. The next router is always one of the adjacent neighbors or the IP address of the local interface for a directly-attached network.
<i>Next Hop Interface</i>	The outgoing interface to use when forwarding traffic to the destination. For a static reject route, the next hop is Null.
<i>Best Route</i>	Indicates whether the route is the preferred route to the network. If the field is blank, a better route to the same network exists in the routing table.

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

Configured Routes

Use the *Configured Route Summary* page to create and display static routes.

To display the page, click **Routing > Router > Configured Routes** in the navigation menu.



Configured Route Summary

Configured Route Summary Fields

Field	Description
<i>Network Address</i>	The IP route prefix for the destination network. This IP address must contain only the network portion of the address and not the host bits. When adding a default route, this field is not available.
<i>Subnet Mask</i>	The IP subnet mask (also known as the network mask or netmask) associated with the network address. The subnet mask defines which portion of an IP address belongs to the network prefix, and which portion belongs to the host identifier. When adding a default route, this field is not available.
<i>Next Hop IP Address</i>	The outgoing router IP address to use when forwarding traffic to the next router (if any) in the path towards the destination. The next router is always one of the adjacent neighbors or the IP address of the local interface for a directly-attached network. When adding a static reject route, this field is not available because the packets are dropped rather than forwarded.
<i>Next Hop Interface</i>	The outgoing interface to use when forwarding traffic to the destination. For a static reject route, the next hop is Null.
<i>Preference</i>	The preference of the route. A lower preference value indicates a more preferred route. When the routing table has more than one route to the same network, the device selects the route with the best (lowest) route preference.

Use the buttons to perform the following tasks:

- To configure a route, click **Add** and configure the settings as described in **“Adding a Static Route” on page 201**. Then, click **Submit** to apply the changes.
- To remove a configured route, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch’s next power cycle, click **System > Configuration Storage > Save**.

Adding a Static Route

In order to create a route, a valid routing interface must exist and the next hop IP Address must be on the same network as the routing interface. To create a route, use the *Routing IP Interface Configuration* page (refer to **“Routing IP Interface Configuration” on page 195**). To see valid next hop IP addresses, use the *Route Table* page (refer to **“Route Table” on page 200**).

Follow these steps to add a static route from the *Configured Routes* page:

1. Click **Add**.

The *Add Route* dialog box appears.

2. Next to *Route Type*, select one of the following:

- **Default** The route the device uses to send a packet if the routing table does not contain a longer matching prefix for the packet's destination. The routing table can contain only one default route.
- **Static** A route that is manually added to the routing table by an administrator.
- **Static Reject** A route where packets that match the route are discarded instead of forwarded. The device might send an ICMP Destination Unreachable message.

3. Configure the remaining available fields in the *Add Route* dialog box.



Note: The selected *Route Type* determines the fields that are available to be configured. Some of the fields listed in the table **“Configured Routes Fields” on page 201** are not available when configuring certain types of routes.

4. Click **Submit** to apply the changes.

The new route is added, and you are returned to the *Configured Routes* page.

Configuring Policy-Based Routing

Policy-based routing (PBR) enhances/modifies existing features in the EdgeSwitch software. These features are route maps and access-control lists. Route maps are part of routing (see [“Router” on page 200](#)) and access control lists are part of QOS (see [“Configuring Access Control Lists” on page 230](#)). As policy-based routing feature utilizes services of both features mentioned above, the EdgeSwitch software with a combination of Routing and QOS packages is required to have PBR functional.

Normally, routers take forwarding decision based on routing tables in order to forward packets to destination addresses. Policy-Based Routing is a feature that enables network administrator to define forwarding behavior based on packet contents. In brief, Policy-Based Routing overrides traditional destination-based routing behavior.

The EdgeSwitch software’s policy-based routing feature match the following packet entities and overrides traditional forwarding behavior accomplished through destination-based routing:

- The size of the packet
- Protocol of the payload
- Source MAC address
- Destination MAC address
- Source IP address
- Destination IP address
- VLAN tag
- Priority

Chapter 6: Managing Device Security

Use the features in the Security folder on the navigation menu to set management security parameters for port, user, and server security. The Security folder contains links to the following features:

- **[“Port Access Control” on page 204](#)**
- **[“RADIUS Settings” on page 218](#)**
- **[“TACACS+ Settings” on page 225](#)**

Port Access Control

In port-based authentication mode, when 802.1X is enabled globally and on the port, successful authentication of any one supplicant attached to the port results in all users being able to use the port without restrictions. At any given time, only one supplicant is allowed to attempt authentication on a port in this mode. Ports in this mode are under bidirectional control. This is the default authentication mode.

The 802.1X network has three components:

- **Authenticators:** Specifies the port that is authenticated before permitting system access.
- **Supplicants:** Specifies host connected to the authenticated port requesting access to the system services.

Authentication Server: Specifies the external server, for example, the RADIUS server that performs the authentication on behalf of the authenticator, and indicates whether the user is authorized to access system services.

The Port Access Control folder contains links to the following pages that allow you to view and configure 802.1X features on the system.

Global Port Access Control Configuration

Use the *Port Access Control Configuration* page to enable or disable port access control on the system.

To display the *Port Access Control Configuration* page, click **Security > Port Access Control > Configuration** in the navigation menu.

Port Access Control Configuration

Port Access Control Configuration Fields

Field	Description
<i>Admin Mode</i>	Specifies whether to <i>Enable</i> or <i>Disable</i> port-based authentication on the switch. The default is <i>Disable</i> .
<i>VLAN Assignment Mode</i>	The administrative mode of RADIUS-based VLAN assignment on the device. When enabled, this feature allows a port to be placed into a particular VLAN based on the result of the authentication or type of 802.1X authentication a client uses when it accesses the device. The authentication server can provide information to the device about which VLAN to assign the supplicant.
<i>Dynamic VLAN Creation Mode</i>	The administrative mode of dynamic VLAN creation on the device. Select <i>Enable</i> to allow the switch to dynamically create a RADIUS-assigned VLAN if it does not already exist in the VLAN database. If RADIUS-assigned VLANs are enabled, the RADIUS server is expected to include the VLAN ID in the 802.1X tunnel attributes of its response message to the device. If dynamic VLAN creation is enabled on the device and the RADIUS-assigned VLAN does not exist, then the assigned VLAN is dynamically created. This implies that the client can connect from any port and can get assigned to the appropriate VLAN. This feature gives flexibility for clients to move around the network without much additional configuration required.
<i>Monitor Mode</i>	The administrative mode of the Monitor Mode feature on the device. Monitor mode is a special mode that can be enabled in conjunction with port-based access control. Monitor mode provides a way for network administrators to identify possible issues with the port-based access control configuration on the device without affecting the network access to the users of the device. It allows network access even in cases where there is a failure to authenticate, but it logs the results of the authentication process for diagnostic purposes. If the device fails to authenticate a client for any reason (for example, RADIUS access reject from the RADIUS server, RADIUS timeout, or the client itself is 802.1X unaware), the client is authenticated and is undisturbed by the failure condition(s). The reasons for failure are logged and buffered into the local logging database for tracking purposes.
<i>EAPOL Flood Mode</i>	The administrative mode of the Extensible Authentication Protocol (EAP) over LAN (EAPOL) flood support on the device. EAPOL Flood Mode can be enabled when <i>Admin Mode</i> and <i>Monitor Mode</i> are disabled.

Use the buttons to perform the following tasks:

- If you change any settings, click **Submit** to apply the new settings to the system.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Port Access Control Port Summary

Use this page to view summary information about the port-based authentication settings for each port.

To display the *Port Access Control Port Summary* page, click **Security > Port Access Control > Port Summary** in the navigation menu.

Interface	PAE Capabilities	Control Mode	Operating Control Mode	PAE State	Backend State	
0/1	Authenticator	Auto	N/A	Initialize	Initialize	[Init] [Refresh]
0/2	Authenticator	Auto	N/A	Initialize	Initialize	[Init] [Refresh]
0/3	Authenticator	Auto	N/A	Initialize	Initialize	[Init] [Refresh]
0/4	Authenticator	Auto	N/A	Initialize	Initialize	[Init] [Refresh]
0/5	Authenticator	Auto	N/A	Initialize	Initialize	[Init] [Refresh]
0/6	Authenticator	Auto	N/A	Initialize	Initialize	[Init] [Refresh]
0/7	Authenticator	Auto	N/A	Initialize	Initialize	[Init] [Refresh]
0/8	Authenticator	Auto	N/A	Initialize	Initialize	[Init] [Refresh]
0/9	Authenticator	Auto	N/A	Initialize	Initialize	[Init] [Refresh]
0/10	Authenticator	Auto	N/A	Initialize	Initialize	[Init] [Refresh]

Port Access Control Port Summary

Port Access Control Port Summary Fields

Field	Description
<i>Interface</i>	The interface associated with the rest of the data in the row.
<i>PAE Capabilities</i>	The Port Access Entity (PAE) role, which is one of the following: Authenticator The port enforces authentication and passes authentication information from a remote supplicant (similar to a client or host) to the authentication server. If the server successfully authenticates the supplicant, the port allows access. Supplicant The port must be granted permission by the authentication server before it can access the remote authenticator port.

Port Access Control Port Summary Fields (Continued)

Field	Description
<i>Control Mode</i>	The port-based access control mode configured on the port, which is one of the following: • Auto The port is unauthorized until a successful authentication exchange has taken place. • Force Unauthorized The port ignores supplicant authentication attempts and does not provide authentication services to the client. • Force Authorized The port sends and receives normal traffic without client port-based authentication. • MAC-Based This mode allows multiple supplicants connected to the same port to each authenticate individually. Each host connected to the port must authenticate separately in order to gain access to the network. The hosts are distinguished by their MAC addresses.
<i>Operating Control Mode</i>	The control mode under which the port is actually operating, which is one of the following: • Auto • Force Unauthorized • Force Authorized • MAC-Based • N/A If the mode is N/A, port-based access control is not applicable to the port. If the port is in detached state it cannot participate in port access control. Additionally, if port-based access control is globally disabled, the status for all ports is N/A.
<i>PAE State</i>	The current state of the authenticator PAE state machine, which is the 802.1X process that controls access to the port. The state can be one of the following: • Initialize • Disconnected • Connecting • Authenticating • Authenticated • Aborting • Held • ForceAuthorized • ForceUnauthorized
<i>Backend State</i>	The current state of the back-end authentication state machine, which is the 802.1X process that controls the interaction between the 802.1X client on the local system and the remote authentication server. The state can be one of the following: • Request • Response • Success • Fail • Timeout • Initialize • Idle
Command buttons (for each interface)	 Click this button to reset the 802.1X state machine on the associated interface to the initialization state. Traffic sent to and from the port is blocked during the authentication process. This button can be clicked only when the port is an authenticator and the operating <i>Control Mode</i> is Auto.  Click this button to force the associated interface to restart the authentication process.

Use the buttons to perform the following tasks:

- To change the port-based access control settings for a port, select the port and click **Edit**. The UI automatically redirects to the *Port Access Control Port Configuration* page for the selected port. For information on using this page, refer to **“Global Port Access Control Configuration” on page 205**.
- To view additional information about the port-based access control settings for a port, select the port with the information to view and click **Details**. You are automatically redirected to the *Port Access Control Port Details* page for the selected port.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Port Access Control Port Configuration

Use the *Port Access Control Port Configuration* page to enable and configure port access control on one or more ports.

To access the *Port Access Control Port Configuration* page, click **Security > Port Access Control > Port Configuration** in the navigation menu.

The screenshot displays the 'Port Access Control Port Configuration' page. At the top, there are tabs for Configuration, Port Summary, Port Configuration (selected), Port Details, Statistics, Client Summary, and Privileges Summary. The main title is 'Port Access Control Port Configuration'. Below this, the 'Interface' is set to '0/1'. The 'PAE Capabilities' section shows 'Authenticator' as the selected option. The 'Authenticator Options' section includes fields for Control Mode (Auto), Quiet Period (80 seconds), Transmit Period (30 seconds), Guest VLAN ID (empty), Guest VLAN Period (60 seconds), Unauthenticated VLAN ID (empty), Supplicant Timeout (30 seconds), Server Timeout (30 seconds), Maximum Requests (2), MAB Mode (unchecked), Re-Authentication Period (Disabled), and Maximum Users (40). The 'Supplicant Options' section includes Control Mode (Auto), User Name (None), Authentication Period (30 seconds), Start Period (30 seconds), Held Period (60 seconds), and Maximum Start Messages (3). At the bottom, there are 'Submit', 'Refresh', and 'Cancel' buttons. A copyright notice for Ubiquiti Networks, Inc. is visible in the footer.

Port Access Control Port Configuration

Port Access Control Port Configuration Fields

Field	Description
<i>Interface</i>	The interface with the settings to view or configure. If you have been redirected to this page, this field is read-only and displays the interface that was selected on the <i>Port Access Control Port Summary</i> page.
<i>PAE Capabilities</i>	The Port Access Entity (PAE) role, which is one of the following: Authenticator The port enforces authentication and passes authentication information from a remote supplicant (client or host) to the authentication server. If the server successfully authenticates the supplicant, the port allows access. Supplicant The port is connected to an authenticator port and must be granted permission by the authentication server before it can send and receive traffic through the remote port. To change the PAE capabilities of a port, click the  button next to the field and select the desired setting from the drop-down box in the <i>Set PAE Capabilities</i> window.
<i>Authenticator Options</i> – The fields in this section can be changed only when the selected port is configured as an authenticator port (that is, the <i>PAE Capabilities</i> field is set to Authenticator).	
<i>Control Mode</i>	The port-based access control mode on the port, which is one of the following: Auto The port is unauthorized until a successful authentication exchange has taken place. Force Unauthorized The port ignores supplicant authentication attempts and does not provide authentication services to the client Force Authorized The port sends and receives normal traffic without client port-based authentication. MAC-Based This mode allows multiple supplicants connected to the same port to each authenticate individually. Each host connected to the port must authenticate separately in order to gain access to the network. The hosts are distinguished by their MAC addresses.
<i>Quiet Period</i>	The number of seconds that the port remains in the quiet state following a failed authentication exchange.
<i>Transmit Period</i>	The value, in seconds, of the timer used by the authenticator state machine on the port to determine when to send an EAPOL EAP Request/Identity frame to the supplicant.
<i>Guest VLAN ID</i>	The VLAN ID of the guest VLAN. The guest VLAN allows the port to provide a distinguished service to unauthenticated users. This feature is a mechanism to allow users to access hosts on the guest VLAN.  Click this button to set the <i>Guest VLAN ID</i>.  Click this button to reset the <i>Guest VLAN ID</i> to the default value.
<i>Guest VLAN Period</i>	The value, in seconds, of the timer used for guest VLAN authentication.
<i>Unauthenticated VLAN ID</i>	The VLAN ID of the unauthenticated VLAN. Hosts that fail the authentication might be denied access to the network or placed on a VLAN created for unauthenticated clients. This VLAN might be configured with limited network access.  Click this button to set the <i>Unauthenticated VLAN ID</i>.  Click this button to reset the <i>Unauthenticated VLAN ID</i> to the default value.
<i>Supplicant Timeout</i>	The amount of time that the port waits for a response before retransmitting an EAP request frame to the client.
<i>Server Timeout</i>	The amount of time the port waits for a response from the authentication server.
<i>Maximum Requests</i>	The maximum number of times that the port sends an EAP request frame (assuming that no response is received) to the client before restarting the authentication process.
<i>MAB Mode</i>	The MAC-based Authentication Bypass (MAB) mode on the port, which can be enabled or disabled.
<i>Re-Authentication Period</i>	The amount of time that clients can be connected to the port without being reauthenticated. If this field is disabled, connected clients are not forced to reauthenticate periodically.  Click this button to set the <i>Re-Authentication Period</i>.  Click this button to reset the <i>Re-Authentication Period</i> to the default value.
<i>Maximum Users</i>	The maximum number of clients supported on the port if the <i>Control Mode</i> on the port is MAC-Based 802.1X authentication.

Port Access Control Port Configuration Fields (Continued)

Field	Description
<i>Supplicant Options</i> – The fields in this section can be changed only when the selected port is configured as a supplicant port (that is, the <i>PAE Capabilities</i> field is set to Supplicant).	
<i>Control Mode</i>	The port-based access control mode on the port, which is one of the following: • Auto The port is in an unauthorized state until a successful authentication exchange has taken place between the supplicant port, the authenticator port, and the authentication server. • Force Unauthorized The port is placed into an unauthorized state and is automatically denied system access. • Force Authorized The port is placed into an authorized state and does not require client port-based authentication to be able to send and receive traffic.
<i>User Name</i>	The name the port uses to identify itself as a supplicant to the authenticator port. The menu includes the users that are configured for system management. When authenticating, the supplicant provides the password associated with the selected <i>User Name</i> .
<i>Authentication Period</i>	The amount of time the supplicant port waits to receive a challenge from the authentication server. If the configured <i>Authentication Period</i> expires, the supplicant retransmits the authentication request until it is authenticated or has sent the number of messages configured in the <i>Maximum Start Messages</i> field.
<i>Start Period</i>	The amount of time the supplicant port waits for a response from the authenticator port after sending a Start packet. If no response is received, the supplicant retransmits the Start packet.
<i>Held Period</i>	The amount of time the supplicant port waits before contacting the authenticator port after an active 802.1X session fails.
<i>Maximum Start Messages</i>	The maximum number of Start packets the supplicant port sends to the authenticator port without receiving a response before it considers the authenticator to be 802.1X-unaware.

Use the buttons to perform the following tasks:

- If you change any settings on this page, click **Submit** to apply the changes.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Port Access Control Port Details

Use this page to view 802.1X information for a specific port.

To access the *Port Access Control Port Details* page, click **Security** > **Port Access Control** > **Port Details** in the navigation menu.

The screenshot displays the 'Port Access Control Port Details' page. At the top, there are tabs for Configuration, Port Summary, Port Configuration, Port Details (selected), Statistics, Client Summary, and Privileges Summary. Below the tabs, the 'Interface' is set to '0/1'. The 'PAE Capabilities' are set to 'Authenticator'. The 'Authenticator Options' section contains a table of settings:

Authenticator Options	
Control Mode	Auto
Quiet Period (Seconds)	60
Transmit Period (Seconds)	30
Guest VLAN ID	0
Guest VLAN Period (Seconds)	90
Unauthenticated VLAN ID	0
Supplicant Timeout (Seconds)	30
Server Timeout (Seconds)	30
Maximum Requests	2
Configured NAS Mode	Disabled
Operational NAS Mode	Disabled
Re-Authentication Period (Seconds)	Disabled
Maximum Users	48

A 'Refresh' button is located at the bottom right of the table. The copyright notice at the bottom right reads '© Copyright 2012-2014 Ubiquiti Networks, Inc.'

Port Access Control Port Details

Port Access Control Port Details Fields

Field	Description
<i>Interface</i>	The interface associated with the rest of the data on the page.
<i>PAE Capabilities</i>	The Port Access Entity (PAE) role, which is one of the following: Authenticator The port enforces authentication and passes authentication information from a remote supplicant (client or host) to the authentication server. If the server successfully authenticates the supplicant, the port allows access. Supplicant The port is connected to an authenticator port and must be granted permission by the authentication server before it can send and receive traffic through the remote port.
Authenticator Options – The fields in this section provide information about the settings that apply to the port when it is configured as an 802.1X authenticator.	
<i>Control Mode</i>	The port-based access control mode on the port, which is one of the following: Auto The port is unauthorized until a successful authentication exchange has taken place. Force Unauthorized The port ignores supplicant authentication attempts and does not provide authentication services to the client. Force Authorized The port sends and receives normal traffic without client port-based authentication. MAC-Based This mode allows multiple supplicants connected to the same port to each authenticate individually. Each host connected to the port must authenticate separately in order to gain access to the network. The hosts are distinguished by their MAC addresses.
<i>Quiet Period</i>	The number of seconds that the port remains in the quiet state following a failed authentication exchange.

Port Access Control Port Details Fields (Continued)

Field	Description
<i>Transmit Period</i>	The value, in seconds, of the timer used by the authenticator state machine on the port to determine when to send an EAPOL EAP Request/Identity frame to the supplicant.
<i>Guest VLAN ID</i>	The VLAN ID for the guest VLAN. The guest VLAN allows the port to provide a distinguished service to unauthenticated users. This feature provides a mechanism to allow users access to hosts on the guest VLAN.
<i>Guest VLAN Period</i>	The value, in seconds, of the timer used for guest VLAN authentication.
<i>Unauthenticated VLAN ID</i>	The VLAN ID of the unauthenticated VLAN. Hosts that fail the authentication might be denied access to the network or placed on a VLAN created for unauthenticated clients. This VLAN might be configured with limited network access.
<i>Supplicant Timeout</i>	The amount of time that the port waits for a response before retransmitting an EAP request frame to the client.
<i>Server Timeout</i>	The amount of time the port waits for a response from the authentication server.
<i>Maximum Requests</i>	The maximum number of times that the port sends an EAP request frame (assuming that no response is received) to the client before restarting the authentication process.
<i>Configured MAB Mode</i>	<i>The configured MAC-based Authentication Bypass (MAB) mode on the port.</i>
<i>Operational MAB Mode</i>	<i>The operational MAC-based Authentication Bypass (MAB) mode on the port.</i>
<i>Re-Authentication Period</i>	The amount of time that clients can be connected to the port without being reauthenticated. If this field is disabled, connected clients are not forced to reauthenticate periodically.
<i>Maximum Users</i>	The maximum number of clients supported on the port if the Control Mode on the port is MAC-based 802.1X authentication.

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

Port Access Control Statistics

Use this page to view information about the Extensible Authentication Protocol over LAN (EAPOL) frames and EAP messages sent and received by the local interfaces. To view additional per-interface EAPOL and EAP message statistics, select the interface with the information to view and click **Details**.

To access the *Port Access Control Statistics* page, click **Security > Port Access Control > Statistics** in the navigation menu.

Interface	PAE Capabilities	EAPOL Frames Received	EAPOL Frames Transmitted	Last EAPOL Frame Version	Last EAPOL Frame Source
0/1	Authenticator	0	0	0	08:00:00:00:00:00
0/2	Authenticator	0	0	0	08:00:00:00:00:00
0/3	Authenticator	0	0	0	08:00:00:00:00:00
0/4	Authenticator	0	0	0	08:00:00:00:00:00
0/5	Authenticator	0	0	0	08:00:00:00:00:00
0/6	Authenticator	0	0	0	08:00:00:00:00:00
0/7	Authenticator	0	0	0	08:00:00:00:00:00
0/8	Authenticator	0	0	0	08:00:00:00:00:00
0/9	Authenticator	0	0	0	08:00:00:00:00:00
0/10	Authenticator	0	0	0	08:00:00:00:00:00

Port Access Control Statistics

Port Access Control Statistics Fields

Field	Description
<i>Interface</i>	The interface associated with the rest of the data in the row. When viewing detailed information for an interface, this field identifies the interface being viewed.
<i>PAE Capabilities</i>	The Port Access Entity (PAE) role, which is one of the following: Authenticator The port enforces authentication and passes authentication information from a remote supplicant (similar to a client or host) to the authentication server. If the server successfully authenticates the supplicant, the port allows access. Supplicant The port must be granted permission by the authentication server before it can access the remote authenticator port.
<i>EAPOL Frames Received</i>	The total number of valid EAPOL frames received on the interface.
<i>EAPOL Frames Transmitted</i>	The total number of EAPOL frames sent by the interface.
<i>Last EAPOL Frame Version</i>	The protocol version number attached to the most recently received EAPOL frame.
<i>Last EAPOL Frame Source</i>	The source MAC address attached to the most recently received EAPOL frame.
<i>Details window fields – Click Details to open a window with additional information about the EAPOL and EAP messages the interface sends and receives. The following information describes the additional fields that appear in the <i>Details</i> window. The fields this window displays depend on whether the interface is configured as an authenticator or supplicant, as noted in the applicable field descriptions.</i>	
<i>EAPOL Start Frames Received</i>	The total number of EAPOL-Start frames received on the interface. EAPOL-Start frames are sent by a supplicant to initiate the 802.1X authentication process when it connects to the interface. This field is displayed only if the interface is configured as an authenticator.
<i>EAPOL Logoff Frames Received</i>	The total number of EAPOL-Logoff frames received on the interface. EAPOL-Logoff frames are sent by a supplicant to indicate that it is disconnecting from the network, and the interface can return to the unauthorized state. This field is displayed only if the interface is configured as an authenticator.

Port Access Control Statistics Fields (Continued)

Field	Description
<i>EAP Response/ID Frames Received</i>	The total number of EAP-Response Identity frames the interface has received. EAP-Response Identity frames are sent by a supplicant to provide user information that is used to for authentication. This field is displayed only if the interface is configured as an authenticator.
<i>EAP Response Frames Received</i>	The total number of EAP-Response frames the interface has received. EAP-Response frames are sent from a supplicant to an authentication server during the authentication process. This field is displayed only if the interface is configured as an authenticator.
<i>EAP Request/ID Frames Transmitted</i>	The total number of EAP-Request Identity frames the interface has sent. EAP-Request Identity frames are sent from an authenticator to a supplicant to request user information that is used to for authentication. This field is displayed only if the interface is configured as an authenticator.
<i>EAPOL Start Frames Transmitted</i>	The total number of EAPOL-Start frames the interface has sent to a remote authenticator. EAPOL-Start frames are sent by a supplicant to initiate the 802.1X authentication process when it connects to the interface. This field is displayed only if the interface is configured as a supplicant.
<i>EAPOL Logoff Frames Transmitted</i>	The total number of EAPOL-Logoff frames the interface has sent to a remote authenticator. EAPOL-Logoff frames are sent by a supplicant to indicate that it is disconnecting from the network, and the interface can return to the unauthorized state. This field is displayed only if the interface is configured as a supplicant.
<i>EAP Response/ID Frames Transmitted</i>	The total number of EAP-Response Identity frames the interface has sent. EAP-Response Identity frames are sent by a supplicant to provide user information that is used to for authentication. This field is displayed only if the interface is configured as a supplicant.
<i>EAP Request/ID Frames Received</i>	The total number of EAP-Request Identity frames the interface has received. EAP-Request Identity frames are sent from an authenticator to a supplicant to request user information that is used to for authentication. This field is displayed only if the interface is configured as a supplicant.
<i>EAP Request Frames Received</i>	The total number of EAP-Request frames the interface has received. EAP-Request frames are sent from the authentication server to the supplicant during the authentication process. This field is displayed only if the interface is configured as a supplicant.
<i>Invalid EAPOL Frames Received</i>	The number of unrecognized EAPOL frames received on the interface.
<i>EAPOL Length Error Frames Received</i>	The number of EAPOL frames with an invalid packet body length received on the interface.
<i>Clear (Button)</i>	Resets all statistics counters to 0 for the selected interface or interfaces.

Use the buttons to perform the following tasks:

- Click **Details** to view additional per-interface EAPOL and EAP message statistics for the selected interface(s).
- Click **Clear** to reset all statistics counters to 0 for the selected interface(s).
- Click **Refresh** to refresh the page with the most current data from the switch.

Port Access Control Client Summary

This page displays information about supplicant devices that are connected to the local authenticator ports. If there are no active 802.1X sessions, the table is empty. To view additional information about a supplicant, select the interface it is connected to and click **Details**.

To access the *Port Access Control Client Summary* page, click **Security** > **Port Access Control** > **Client Summary** in the navigation menu.



Port Access Control Client Summary

Port Access Control Client Summary Fields

Field	Description
<i>Interface</i>	The local interface associated with the rest of the data in the row. When viewing detailed information for an interface, this field identifies the interface being viewed.
<i>Logical Interface</i>	The logical port number associated with the supplicant that is connected to the port.
<i>User Name</i>	The name the client uses to identify itself as a supplicant to the authentication server.
<i>Supplicant MAC Address</i>	The MAC address of the supplicant that is connected to the port.
<i>Session Time</i>	The amount of time that has passed since the connected supplicant was granted access to the network through the authenticator port.
<i>Filter ID</i>	The policy filter ID assigned by the authenticator to the supplicant device.
<i>VLAN ID</i>	The ID of the VLAN the supplicant was placed in as a result of the authentication process.
Details Window Fields	
After you click Details , a window opens and displays the following additional information about the client:	
<i>Session Timeout</i>	The reauthentication timeout period set by the RADIUS server to the supplicant device.
<i>Session Termination Action</i>	The termination action set by the RADIUS server that indicates the action that will take place once the supplicant reaches the session timeout value.


Use the buttons to perform the following tasks:

- Click **Details** to view additional information for the selected client(s), as shown in the table above.
- Click **Refresh** to refresh the page with the most current data from the switch.

Port Access Control Privileges Summary

Use this page to grant or deny port access to users configured on the system. To change the access control privileges for one or more ports, select each interface to configure and click **Edit**. The same settings are applied to all selected interfaces.

To access the *Port Access Control Privileges Summary* page, click **Security > Port Access Control > Privileges Summary** in the navigation menu.



Port Access Control Privileges Summary

Port Access Control Privileges Summary Fields

Field	Description
<i>Interface</i>	The local interface associated with the rest of the data in the row. When configuring access information for one or more interfaces, this field identifies each interface being configured.
<i>Users</i>	The users that are allowed access to the system through the associated port. When configuring user access for a port, the <i>Available Users</i> field lists the users configured on the system who are denied access to the port, while the <i>Selected Users</i> field lists users who are allowed access to the port. To move a user from one field to the other, click the user to move (or press and hold CTRL to select multiple users) and click or .

Use the buttons to perform the following tasks:


- To change the access control privileges for one or more ports, select the interface(s) to configure and click **Edit**. Configure the settings as needed and click **Submit** to apply the changes. The same settings are applied to all selected interfaces.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Port Access Control History Log Summary

Use this page to grant or deny port access to users configured on the system. To change the access control privileges for one or more ports, select each interface to configure and click Edit. The same settings are applied to all selected interfaces.

To access the *Port Access Control History Log Summary* page, click **Security > Port Access Control > History Log Summary** in the navigation menu.



Port Access Control History Log Summary

Port Access Control History Log Summary Fields

Field	Description
<i>Interface</i>	The interface associated with the rest of the data in the row. Only interfaces that have entries in the log history are listed.
<i>Time Stamp</i>	The absolute time when the authentication event took place.
<i>VLAN Assigned</i>	The ID of the VLAN the supplicant was placed in as a result of the authentication process.
<i>VLAN Assigned Reason</i>	The reason why the authenticator placed the supplicant in the VLAN. Possible values are: RADIUS Unauth Default Not Assigned
<i>Supp MAC Address</i>	The MAC address of the supplicant that is connected to the port.
<i>Filter Name</i>	The policy filter ID assigned by the authenticator to the supplicant device.
<i>Auth Status</i>	The authentication status of the client or port.
<i>Reason</i>	The reason for the successful or unsuccessful authentication.

Use the buttons to perform the following tasks:

- To clear the history log, click **Clear History**.
- Click **Refresh** to refresh the page with the most current data from the switch.

RADIUS Settings

Remote Authorization Dial-In User Service (RADIUS) servers provide additional security for networks. The RADIUS server maintains a user database, which contains per-user authentication information. RADIUS servers provide a centralized authentication method for:

- Telnet Access
- Web Access
- Console to Switch Access
- Port Access Control (802.1X)

The RADIUS folder contains links to pages that help you view and configure system RADIUS settings.

RADIUS Configuration

Use the *RADIUS Configuration* page to view and configure various settings for the RADIUS servers configured on the system. To access the page, click **Security > RADIUS > Configuration** in the navigation menu.

RADIUS Configuration

RADIUS Configuration Fields

Field	Description
<i>Max Number of Retransmits</i>	The maximum number of times the RADIUS client on the device will retransmit a request packet to a configured RADIUS server after a response is not received. If multiple RADIUS servers are configured, the max retransmit value will be exhausted on the first server before the next server is attempted. A retransmit will not occur until the configured timeout value on that server has passed without a response from the RADIUS server. Therefore, the maximum delay in receiving a response from the RADIUS server equals the sum of (retransmit × timeout) for all configured servers. If the RADIUS request was generated by a user login attempt, all user interfaces will be blocked until the RADIUS application returns a response.
<i>Timeout Duration</i>	The number of seconds the RADIUS client waits for a response from the RADIUS server. Consideration to maximum delay time should be given when configuring RADIUS timeout and RADIUS max retransmit values.
<i>Accounting Mode</i>	Used to <i>Enable</i> or <i>Disable</i> the RADIUS accounting mode on the device.
<i>NAS-IP Address</i>	The network access server (NAS) IP address for the RADIUS server. Click this button to specify the NAS IP address. The address should be unique to the NAS within the scope of the RADIUS server. The NAS IP address is used only in Access-Request packets. Click this button to reset the <i>NAS IP Address</i> to the default value.

Use the buttons at the bottom of the page to perform the following actions:

- If you make changes to the page, click **Submit** to apply the changes to the system.
- Click **Refresh** to update the page with the most current information.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

RADIUS Named Server Status

The *RADIUS Named Server Status* page shows summary information about the RADIUS servers configured on the system.

To access the page, click **Security > RADIUS > Named Server** in the navigation menu.



RADIUS Named Server Status

RADIUS Server Status Fields

Field	Description
<i>Current</i>	An asterisk (*) in the column Indicates that the server is the current server for the authentication server group. If no asterisk is present, the server is a backup server. If more than one RADIUS server is configured with the same name, the switch selects one of the servers to be the current server from the group of servers with the same name. When the switch sends a RADIUS request to the named server, the request is directed to the server selected as the current server. Initially the primary server is selected as the current server. If the primary server fails, one of the other servers becomes the current server.
<i>IP Address/Host Address</i>	Shows the IP address of the RADIUS server.
<i>Server Name</i>	Shows the RADIUS server name. Multiple RADIUS servers can have the same name. In this case, RADIUS clients can use RADIUS servers with the same name as backups for each other.
<i>Port Number</i>	Identifies the authentication port the server uses to verify the RADIUS server authentication. The port is a UDP port.
<i>Server Type</i>	Shows whether the server is a <i>Primary</i> or <i>Secondary</i> server.
<i>Secret Configured</i>	Indicates whether the shared secret for this server has been configured.
<i>Message Authenticator</i>	Shows whether the message authenticator attribute for the selected server is enabled or disabled.

Use the buttons to perform the following tasks:

- To add a RADIUS authentication server to the list of servers the RADIUS client can contact, click **Add**, configure the fields, and click **Submit** to apply the changes.
- To change the settings for a configured RADIUS server, select the entry to edit, click **Edit**, configure the fields, and click **Submit** to apply the changes.
- To remove a RADIUS server from the list, select the server, click **Remove**, and confirm the deletion.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

RADIUS Server Statistics

Use the *RADIUS Server Statistics* page to view statistical information for each RADIUS server configured on the system.

To access the *RADIUS Server Statistics* page, click **Security** > **RADIUS** > **Statistics** in the navigation menu.



RADIUS Server Statistics

RADIUS Server Statistics Fields

Field	Description
<i>IP Address/Host Name</i>	The IP address or host name of the RADIUS server associated with the rest of the data in the row. When viewing the detailed statistics for a RADIUS server, this field identifies the RADIUS server.
<i>Round Trip Time</i>	The time interval, in hundredths of a second, between the most recent Access-Reply/Access-Challenge and the Access-Request that matched it from the RADIUS authentication server.
<i>Access Requests</i>	The number of RADIUS Access-Request packets sent to the server. This number does not include retransmissions.
<i>Access Rejects</i>	The number of RADIUS Access-Reject packets, including both valid and invalid packets, that were received from the server.
<i>Pending Requests</i>	The number of RADIUS Access-Request packets destined for the server that have not yet timed out or received a response.
<i>Timeouts</i>	The number of times a response was not received from the server within the configured timeout value.
<i>Packets Dropped</i>	The number of RADIUS packets received from the server on the authentication port and dropped for some other reason.
When you click Details , the <i>RADIUS Server Detailed Statistics</i> window displays the following additional statistics about the number and type of messages sent between the selected RADIUS server and the RADIUS client on the switch:	
<i>Access Retransmissions</i>	The number of RADIUS Access-Request packets that had to be retransmitted to the server because the initial Access-Request packet failed to be successfully delivered.
<i>Access Accepts</i>	The number of RADIUS Access-Accept packets, including both valid and invalid packets, that were received from the server.
<i>Access Challenges</i>	The number of RADIUS Access-Challenge packets, including both valid and invalid packets, that were received from the server.
<i>Malformed Access Responses</i>	The number of malformed RADIUS Access-Response packets received from the server. Malformed packets include packets with an invalid length. Bad authenticators, signature attributes, and unknown types are not included as malformed access responses.
<i>Bad Authenticators</i>	The number of RADIUS Access-Response packets containing invalid authenticators or signature attributes received from the server.
<i>Unknown Types</i>	The number of RADIUS packets of unknown type which were received from the server on the authentication port.

Use the buttons to perform the following tasks:

- To display additional statistics information listed in the table above, click **Details**.
- Click **Refresh** to refresh the page with the most current data from the switch.

RADIUS Accounting Server Status

The *RADIUS Accounting Server Status* page shows summary information about the accounting servers configured on the system.



RADIUS Accounting Server Status

RADIUS Accounting Server Status Fields

Field	Description
<i>IP Address/Host Name</i>	The IP address or host name of the RADIUS accounting server. Host names must be resolvable by DNS and are composed of a series of labels separated by dots.
<i>Server Name</i>	The name of the RADIUS accounting server. RADIUS servers that are configured with the same name are members of the same named RADIUS server group. RADIUS accounting servers in the same group serve as backups for each other.
<i>Port Number</i>	The UDP port on the RADIUS accounting server to which the local RADIUS client sends request packets.
<i>Secret Configured</i>	Indicates whether the shared secret for this server has been configured.
<i>Secret</i>	The shared secret text string used for authenticating and encrypting all RADIUS communications between the RADIUS client on the device and the RADIUS accounting server. The secret specified in this field must match the shared secret configured on the RADIUS accounting server.

Use the buttons to perform the following tasks:

- To add a RADIUS accounting server to the list of servers the RADIUS client can contact, click **Add**, and click **Submit** to apply the changes.
- To change the settings for a configured RADIUS accounting server, select the entry to modify, click **Edit**, configure the settings as needed (you cannot edit the *IP Address/Host Name*), and click **Submit** to apply the changes.
- To remove a configured RADIUS accounting server from the list, select the entry to delete and click **Remove**. You must confirm the action before the entry is deleted.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

RADIUS Accounting Server Statistics

Use the *RADIUS Accounting Server Statistics* page to view statistical information for each RADIUS server configured on the system.

To access the *RADIUS Accounting Server Statistics* page, click **Security** > **RADIUS** > **Accounting Statistics** in the navigation menu.



RADIUS Accounting Server Statistics

RADIUS Accounting Server Statistics Fields

Field	Description
<i>IP Address/Host Name</i>	The IP address or host name of the RADIUS accounting server associated with the rest of the data in the row. When viewing the detailed statistics for a RADIUS accounting server, this field identifies the server.
<i>Round Trip Time</i>	Displays the time interval, in hundredths of a second, between the most recent Accounting-Response and the Accounting-Request that matched it from this RADIUS accounting server.
<i>Accounting Requests</i>	The number of RADIUS Accounting-Request packets sent to this server. This number does not include retransmissions.
<i>Pending Requests</i>	The number of RADIUS Accounting-Request packets destined for the server that have not yet timed out or received a response.
<i>Timeouts</i>	The number of times a response was not received from the server within the configured timeout value.
<i>Packets Dropped</i>	The number of RADIUS packets received from the server on the accounting port and dropped for some other reason.
<i>Accounting Retransmissions</i>	The number of RADIUS Accounting-Request packets retransmitted to the server.
<i>Accounting Responses</i>	The number of RADIUS packets received on the accounting port from the server.
<i>Timeouts</i>	The number of accounting timeouts to this server.
<i>Malformed Access Responses</i>	The number of malformed RADIUS Accounting-Response packets received from the server. Malformed packets include packets with an invalid length. Bad authenticators and unknown types are not included as malformed accounting responses.
<i>Bad Authenticators</i>	The number of RADIUS Accounting-Response packets that contained invalid authenticators received from the accounting server.
<i>Unknown Types</i>	The number of RADIUS packets of unknown type which were received from the server on the accounting port.

Use the buttons to perform the following tasks:

- Click **Details** to display additional statistics information about the number and type of messages sent between the selected RADIUS server and the RADIUS client on the device.
- Click **Refresh** to refresh the page with the most current data from the switch.

RADIUS Clear Statistics

Use the *RADIUS Clear Statistics* page to reset all RADIUS authentication and accounting statistics to zero.

To access the *RADIUS Clear Statistics* page, click **Security > RADIUS > Clear Statistics** in the navigation menu.



RADIUS Clear Statistics

Click **Reset** to clear all statistics for the RADIUS authentication and accounting server. After you confirm the action, the statistics on both the *RADIUS Server Statistics* and *RADIUS Accounting Server Statistics* pages are reset.

RADIUS Source Interface Configuration

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

To access the *RADIUS Source Interface Configuration* page, click **Security > RADIUS > Source Interface Configuration** in the navigation menu.

RADIUS Source Interface Configuration

RADIUS Source Interface Configuration Fields

Field	Description
Type	The type of interface to use as the source interface: None The primary IP address of the originating (outbound) interface is used as the source address. Interface The primary IP address of a physical port is used as the source address. VLAN The primary IP address of a VLAN routing interface is used as the source address.
Interface	When the selected Type is <i>Interface</i> , select the physical port to use as the source interface.
VLAN ID	When the selected Type is <i>VLAN</i> , select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.

Use the buttons to perform the following tasks:

- If you change any settings on this page, click **Submit** to apply the changes.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

TACACS+ Settings

The TACACS+ submenu allows you to access the pages used to view and modify the TACACS+ configuration.

TACACS+ Configuration

To access the *TACACS+ Configuration* page, click **Security** > **TACACS+** > **Configuration** in the navigation menu.

TACACS+ Configuration

TACACS+ Configuration Fields

Field	Description
<i>Key String</i>	Specifies the authentication and encryption key for TACACS+ communications between the device and the TACACS+ server. The key must match the key configured on the TACACS+ server. Configure Click this button to configure the field. Reset Click this button to reset the field to the default value.
<i>Connection Timeout</i>	The maximum number of seconds allowed to establish a TCP connection between the device and the TACACS+ server.

Use the buttons to perform the following tasks:

- If you change any settings on this page, click **Submit** to apply the changes.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

TACACS+ Server Summary

Use the *TACACS+ Server Summary* page to view and configure information about the TACACS+ Server(s).

To access the page, click **Security > TACACS+ > Server Summary** in the navigation menu.



TACACS+ Server Summary

TACACS+ Server Summary Fields

Field	Description
<i>Server</i>	Specifies the TACACS+ Server IP address or Hostname.
<i>Priority</i>	Specifies the order in which the TACACS+ servers are used.
<i>Port</i>	Specifies the authentication port.
<i>Connection Timeout</i>	The amount of time that passes before the connection between the device and the TACACS+ server times out.
<i>Add TACACS+ Server dialog box</i> – When you click Add , this dialog box appears, allowing you to add a TACACS+ server by configuring the preceding fields, as well as the additional field below:	
<i>Key String</i>	Specifies the authentication and encryption key for TACACS+ communications between the device and the TACACS+ server. The key must match the encryption used on the TACACS+ server.

Use the buttons to perform the following tasks:

- To add a TACACS+ server to the list of servers the TACACS+ client can contact, click **Add** (this button is disabled if the maximum number of servers has already been added). Configure the fields shown in the table below and click **Submit** to apply the changes.
- To edit a configured TACACS+ server from the list, select the entry and click **Edit**. Configure the settings as needed and click **Submit** to apply the changes.
- To remove a configured TACACS+ server from the list, select the entry to delete and click **Remove**. You must confirm the action before the entry is deleted.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

TACACS+ Server Configuration

Use this page to view and configure information about the TACACS+ server(s) that have been configured using the *TACACS+ Server Summary* page (see ***“TACACS+ Server Summary” on page 226***).

To access the *TACACS+ Server Configuration* page, click **Security > TACACS+ > Server Configuration** in the navigation menu.

TACACS+ Server Configuration

TACACS+ Server Configuration Fields

Field	Description
<i>Server</i>	Specifies the TACACS+ Server IP address or Hostname.
<i>Priority</i>	Specifies the order in which the TACACS+ servers are used.
<i>Port</i>	Specifies the authentication port.
<i>Key String</i>	Specifies the authentication and encryption key for TACACS+ communications between the device and the TACACS+ server. The key must match the encryption used on the TACACS+ server.
<i>Connection Timeout</i>	The amount of time that passes before the connection between the device and the TACACS+ server timeout.

Use the buttons to perform the following tasks:

- If you change any fields on this page, click **Submit** to apply the changes.
- To remove a configured TACACS+ server, select it from the table, click **Remove**, and confirm the deletion.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

TACACS+ Source Interface Configuration

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

To access the *TACACS+ Source Interface Configuration* page, click **Security** > **TACACS+** > **Source Interface Configuration** in the navigation menu.

TACACS+ Source Interface Configuration

TACACS+ Source Interface Configuration Fields

Field	Description
<i>Type</i>	The type of interface to use as the source interface: None The primary IP address of the originating (outbound) interface is used as the source address. Interface The primary IP address of a physical port is used as the source address. VLAN The primary IP address of a VLAN routing interface is used as the source address.
<i>Interface</i>	When the selected <i>Type</i> is <i>Interface</i> , select the physical port to use as the source interface.
<i>VLAN ID</i>	When the selected <i>Type</i> is <i>VLAN</i> , select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.

Use the buttons to perform the following tasks:

- If you change any settings on this page, click **Submit** to apply the changes.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

Chapter 7: Configuring Quality of Service

This section gives an overview of Quality of Service (QoS) and explains the QoS features available from the Quality of Service navigation menu. This section contains the following subsections:

- **[“Configuring Access Control Lists” on page 230](#)**
- **[“Configuring Auto VoIP” on page 238](#)**
- **[“Configuring Class of Service” on page 243](#)**

In a typical switch, each physical port consists of one or more queues for transmitting packets on the attached network. Multiple queues per port are often provided to give preference to certain packets over others based on user-defined criteria. When a packet is queued for transmission in a port, the rate at which it is serviced depends on how the queue is configured and possibly the amount of traffic present in the other queues of the port. If a delay is necessary, packets get held in the queue until the scheduler authorizes the queue for transmission. As queues become full, packets have no place to be held for transmission and get dropped by the switch.

QoS is a means of providing consistent, predictable data delivery by distinguishing between packets that have strict timing requirements from those that are more tolerant of delay. Packets with strict timing requirements are given “special treatment” in a QoS capable network. With this in mind, all elements of the network must be QoS-capable. The presence of at least one node which is not QoS-capable creates a deficiency in the network path and the performance of the entire packet flow is compromised.

Configuring Access Control Lists

Access Control Lists (ACLs) ensure that only authorized users have access to specific resources while blocking off any unwarranted attempts to reach network resources. ACLs are used to provide traffic flow control, restrict contents of routing updates, decide which types of traffic are forwarded or blocked, and above all provide security for the network. The EdgeSwitch software supports IPv4 and MAC ACLs. The total number of MAC and IP ACLs supported by the EdgeSwitch software is platform-specific.

You first create an IPv4-based or MAC-based rule and assign a unique ACL ID. Then, you define the rules, which can identify protocols, source and destination IP and MAC addresses, and other packet-matching criteria. Finally, you use the ID number to assign the ACL to a port or to a VLAN interface.

IP Access Control Lists

IP access control lists (ACL) allow network managers to define classification actions and rules for specific ports. ACLs are composed of access control entries (ACE), or rules, that consist of the filters that determine traffic classifications. The total number of rules that can be defined for each ACL is platform-specific. These rules are matched sequentially against a packet. When a packet meets the match criteria of a rule, the specified rule action (Permit/Deny) is taken, including dropping the packet or disabling the port, and the additional rules are not checked for a match. For example, a network administrator defines an ACL rule that says port number 20 can receive TCP packets. However, if a UDP packet is received the packet is dropped.

The IP Access Control List folder contains links to UI pages that allow you to configure and view IP ACLs.

To configure an IP ACL:

1. Use the *IP ACL Configuration* page to define the IP ACL type and assign an ID to it.
2. Use the *Access Control List Interface Summary* page to create rules for the ACL.
3. Use the *Access Control List Configuration* page to view the configuration.

Access Control List Summary

Use the *Access Control List Summary* page to add or remove IP-based ACLs. On this menu the interfaces to which an IP ACL applies must be specified, as well as whether it applies to inbound or outbound traffic.

To display the page, click **QoS > Access Control Lists > Summary** in the navigation menu.



Access Control List Summary

Access Control List Summary Fields

Field	Description
<i>ACL Identifier</i>	The name or number that identifies the ACL. The permitted identifier depends on the ACL type. Standard and Extended IPv4 ACLs use numbers within a set range, and Named IPv4 and MAC ACLs use alphanumeric characters.
<i>ACL Type</i>	The type of ACL. The ACL type determines the criteria that can be used to match packets. The type also determines which attributes can be applied to matching traffic. IPv4 ACLs classify Layer-3 and Layer-4 IPv4 traffic, IPv6 ACLs classify Layer-3 and Layer-4 IPv6 traffic, and MAC ACLs classify Layer-2 traffic. The ACL types are as follows: IPv4 Standard Match criteria is based on the source address of IPv4 packets. IPv4 Extended Match criteria can be based on the source and destination addresses, source and destination Layer-4 ports, and protocol type of IPv4 packets. IPv4 Named Match criteria is the same as IPv4 Extended ACLs, but the ACL ID can be an alphanumeric name instead of a number. IPv6 Named Match criteria can be based on information including the source and destination IPv6 addresses, source and destination Layer-4 ports, and protocol type within IPv6 packets. Extended MAC Match criteria can be based on the source and destination MAC addresses, 802.1p user priority, VLAN ID, and EtherType value within Ethernet frames.
<i>Rules Used</i>	The number of rules currently configured for the ACL.
<i>Direction</i>	Indicates whether the packet is checked against the rules in an ACL when it is received on an interface (<i>Inbound</i>) or after it has been received, routed, and is ready to exit an interface (<i>Outbound</i>).
<i>Interface</i>	The interface(s) to which the ACL has been applied.
<i>VLAN</i>	Each VLAN to which the ACL has been applied.

Use the buttons at the bottom of the page to perform the following tasks:

- To add an ACL, click **Add**, configure the ACL type and ID, and click **Submit** to apply the changes.
- To configure rules for an ACL, select the ACL and click **Edit**. Configure the fields on the *Access Control List Configuration* page for the selected ACL (see [“Access Control List Configuration” on page 232](#)), and click **Submit** to apply the changes.
- To remove one or more configured ACLs, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Access Control List Configuration

Use the *Access Control List Configuration* page to configure rules for existing ACLs on the system and to view summary information about rules that have been added to an ACL. Each ACL rule is configured to match one or more aspects of traffic on the network. When a packet matches all the conditions in a rule, it is handled according to the rule's configured action (permit or deny) and attributes. Each ACL can have multiple rules, but the final rule for every ACL is an implicit deny all rule.

To display the page, click **QoS > Access Control Lists > Configuration** in the navigation menu.



Access Control List Configuration

Access Control List Configuration Fields

Field	Description
<i>ACL Identifier</i>	This drop-down list contains the ID for each ACL on the system. To add or remove a rule, first select the associated ACL's ID from this list. For ACLs with alphanumeric names, click to change the ACL ID. The ID of a Named IPv4 ACL must begin with a letter, and not a number. The ACL identifier for IPv4 Standard and IPv4 Extended ACLs cannot be changed.
<i>Rule</i>	The number that identifies the rule. A number is automatically assigned to a rule when it is created. Rules are added in the order they are created and cannot be renumbered. Packets are checked against the rule criteria in order, from the lowest-numbered rule to the highest. When the packet matches the criteria in a rule, it is handled according to the rule action and attributes. If no rule matches a packet, the packet is discarded based on the implicit deny all rule, which is the final rule in every ACL.
<i>ACL Type</i>	The type of ACL. The ACL type determines the criteria that can be used to match packets. The type also determines which attributes can be applied to matching traffic.
<i>Status</i>	Indicates whether the ACL is active. If the ACL is a time-based ACL that includes a time range, the ACL is active only during the periods specified within the time range. If an ACL does not include a time range, the status is always active.
<i>Action</i>	The action to take when a packet or frame matches the criteria in the rule: • Permit The packet or frame is forwarded. • Deny The packet or frame is dropped. Note: When configuring ACL rules in the <i>Add Access Control List Rule</i> window, the selected action determines which fields can be configured. Not all fields are available for both Permit and Deny actions.
<i>Match Conditions</i>	The criteria used to determine whether a packet or frame matches the ACL rule.
<i>Rule Attributes</i>	Each action, beyond the basic Permit and Deny actions, to perform on the traffic that matches the rule.

Access Control List Configuration Fields (Continued)

Field	Description
Add IPv4 ACL Rule window fields – After you click Add Rule, this window opens, allowing you to add a rule to the ACL selected in the <i>ACL Identifier</i> field. The fields available in the window depend on the <i>ACL Type</i>. The following information describes the fields in this window. The Match Criteria tables that apply to IPv4 ACLs, IPv6 ACLs, and MAC ACLs are described separately. <i>Match Criteria (IPv4 ACLs)</i> – Fields in this section specify the criteria to use to determine whether an IP packet matches the rule. Note: The fields described below apply to IPv4 Standard, IPv4 Extended, and IPv4 Named ACLs, except those marked with an asterisk (*) which apply to IPv4 Extended and IPv4 Named ACLs only.	
<i>Every</i>	When this option is selected, all packets will match the rule and will be either permitted or denied. This option is exclusive to all other match criteria – if <i>Every</i> is selected, no other match criteria can be configured. To configure specific match criteria, this option must be cleared.
<i>Protocol*</i>	The IANA-assigned protocol number to match within the IP packet. You can also specify one of the following keywords: <i>EIGRP, GRE, ICMP, IGMP, IP, IPINIP, OSPF, PIM, TCP, or UDP</i> .
<i>Fragments*</i>	IP ACL rule to match on fragmented IP packets.
<i>Source IP Address / Wildcard Mask</i>	The source port IP address in the packet and source IP wildcard mask (in the second field) to compare to the IP address in a packet header. Wild card masks determines which bits in the IP address are used and which bits are ignored. A wild card mask of <i>255.255.255.255</i> indicates that no bit is important. A wildcard of <i>0.0.0.0</i> indicates that all of the bits are important. Wildcard masking for ACLs operates differently from a subnet mask. A wildcard mask is in essence the inverse of a subnet mask. With a subnet mask, the mask has ones (1's) in the bit positions that are used for the network address, and has zeros (0's) for the bit positions that are not used. In contrast, a wildcard mask has (0's) in a bit position that must be checked. A '1' in a bit position of the ACL mask indicates the corresponding bit can be ignored. This field is required when you configure a source IP address.
<i>Source L4 Port*</i>	The TCP/UDP source port to match in the packet header. The Source L4 Port and Destination L4 port are configurable only if protocol is either <i>TCP</i> or <i>UDP</i> . <i>Equal to, Not Equal to, Greater than, and Less than</i> options are available. For TCP protocol: <i>BGP, Domain, Echo, FTP, FTP-Data, HTTP, SMTP, Telnet, WWW, POP2, or POP3</i> . For UDP protocol: <i>Domain, Echo, NTP, RIP, SNMP, TFTP, Time, or WHO</i> .
<i>Destination IP Address / Wildcard Mask</i>	The destination port IP address in the packet and destination IP wildcard mask (in the second field) to compare to the IP address in a packet header. Wild card masks determines which bits in the IP address are used and which bits are ignored. A wild card mask of <i>255.255.255.255</i> indicates that no bit is important. A wildcard of <i>0.0.0.0</i> indicates that all of the bits are important. Wildcard masking for ACLs operates differently from a subnet mask. A wildcard mask is in essence the inverse of a subnet mask. With a subnet mask, the mask has ones (1's) in the bit positions that are used for the network address, and has zeros (0's) for the bit positions that are not used. In contrast, a wildcard mask has (0's) in a bit position that must be checked. A 1 in a bit position of the ACL mask indicates the corresponding bit can be ignored. This field is required when you configure a destination IP address.
<i>Destination L4 Port*</i>	The TCP/UDP destination port to match in the packet header. The Source L4 Port and Destination L4 port are configurable only if protocol is either <i>TCP</i> or <i>UDP</i> . <i>Equal to, Not Equal to, Greater than, and Less than</i> options are available. For TCP protocol: <i>BGP, Domain, Echo, FTP, FTP-Data, HTTP, SMTP, Telnet, WWW, POP2, or POP3</i> . For UDP protocol: <i>Domain, Echo, NTP, RIP, SNMP, TFTP, Time, or WHO</i> .
<i>IGMP Type*</i>	IP ACL rule to match on the specified IGMP message type. Available only if the protocol is IGMP.
<i>ICMP Type *</i>	IP ACL rule to match on the specified ICMP message type. Available only if the protocol is ICMP.
<i>ICMP Code*</i>	IP ACL rule to match on the specified ICMP message code. Available only if the protocol is ICMP.
<i>ICMP Message*</i>	IP ACL rule to match on the ICMP message type and code. Available only if the protocol is ICMP. Specify one of the following supported ICMP messages: <i>Echo, Echo-Reply, Host-Redirect, Mobile-Redirect, Net-Redirect, Net-Unreachable, Redirect, Packet-Too-Big, Port-Unreachable, Source-Quench, Router-Solicitation, Router-Advertisement, Time-Exceeded, TTL-Exceeded, and Unreachable</i> .
<i>TCP Flags*</i>	IP ACL rule to match on the TCP flags. Available only if the protocol is TCP. When a + flag is specified, a match occurs if the flag is set in the TCP header. When a - flag is specified, a match occurs if the flag is not set in the TCP header. When Established is specified, a match occurs if either RST or ACK bits are set in the TCP header.

Access Control List Configuration Fields (Continued)

Field	Description
<i>Service Type*</i>	The service type to match in the IP header. The available options are alternate ways to specify a match condition for the same <i>Service Type</i> field in the IP header, but each service type uses a different user notation. After you select the service type, specify the value for the service type in the appropriate field. Only the field associated with the selected service type can be configured. The services types are: • IP DSCP Matches the packet IP DiffServ Code Point (DSCP) value to the rule. The DSCP value is defined as the high-order six bits of the Service Type octet in the IP header. • IP Precedence Matches the IP Precedence value to the rule. The IP Precedence field in a packet is defined as the high-order three bits of the Service Type octet in the IP header. • IP TOS Bits Matches on the Type of Service (TOS) bits in the IP header. The IP TOS field in a packet is defined as all eight bits of the Service Type octet in the IP header. For example, to check for an IP TOS value having bits 7 and 5 set and bit 1 clear, where bit 7 is most significant, use a TOS Bits value of 0xA0 and a TOS Mask of 0xFF. • TOS Bits Requires the bits in a packet's TOS field to match the two-digit hexadecimal number entered in this field. • TOS Mask The bit positions that are used for comparison against the IP TOS field in a packet.
<i>Time Range Name</i>	The name of the time range that will impose a time limit on the ACL rule. If a time range with the specified name does not exist, and the ACL containing this rule is associated with an interface, the ACL rule is applied immediately. If a time range with specified name exists, and the ACL containing this ACL rule is associated with an interface, the ACL rule is applied when the time-range with specified name becomes active. The ACL rule is removed when the time-range with specified name becomes inactive.
<i>Committed Rate / Burst Size</i>	The allowed transmission rate for packets on the interface (Committed Rate), and the number of bytes allowed in a temporary traffic burst (Burst Rate).
<i>Match Criteria (IPv6 ACLs)</i> – The fields in this section specify the criteria to use to determine whether an IP packet matches the rule. The fields described below apply to IPv6 ACLs.	
<i>Every</i>	When this option is selected, all packets will match the rule and will be either permitted or denied. This option is exclusive to all other match criteria – if <i>Every</i> is selected, no other match criteria can be configured. To configure specific match criteria, this option must be cleared.
<i>Protocol</i>	The IANA-assigned protocol number to match within the IP packet. You can also specify one of the following keywords: <i>ICMP</i> , <i>IGMP</i> , <i>TCP</i> , <i>UDP</i> , <i>ICMPv6</i> , or <i>IP</i> .
<i>Fragments</i>	IPv6 ACL rule to match on fragmented IP packets.
<i>Source Prefix / Prefix Length</i>	The IPv6 prefix combined with IPv6 prefix length of the network or host sending the packet.
<i>Source L4 Port</i>	The TCP/UDP source port to match in the packet header. Select one of the following options: <i>Equal</i> , <i>Not Equal</i> , <i>Less Than</i> , <i>Greater Than</i> , or <i>Range</i> , and specify the port number or keyword. TCP port keywords include <i>BGP</i> , <i>Domain</i> , <i>Echo</i> , <i>FTP</i> , <i>FTP Data</i> , <i>HTTP</i> , <i>SMTP</i> , <i>Telnet</i> , <i>WWW</i> , <i>POP2</i> , and <i>POP3</i> . UDP port keywords include <i>Domain</i> , <i>Echo</i> , <i>NTP</i> , <i>RIP</i> , <i>SNMP</i> , <i>TFTP</i> , <i>TIME</i> , and <i>WHO</i> .
<i>Destination Prefix / Prefix Length</i>	The IPv6 prefix combined with the IPv6 prefix length to be compared to a packet's destination IPv6 address as a match criteria for the IPv6 ACL rule. To indicate a destination host, specify an IPv6 prefix length of 128.
<i>Destination L4 Port</i>	The TCP/UDP destination port to match in the packet header. Select one of the following options: <i>Equal</i> , <i>Not Equal</i> , <i>Less Than</i> , <i>Greater Than</i> , or <i>Range</i> , and specify the port number or keyword. TCP port keywords include <i>BGP</i> , <i>Domain</i> , <i>Echo</i> , <i>FTP</i> , <i>FTP Data</i> , <i>HTTP</i> , <i>SMTP</i> , <i>Telnet</i> , <i>WWW</i> , <i>POP2</i> , and <i>POP3</i> . UDP port keywords include <i>Domain</i> , <i>Echo</i> , <i>NTP</i> , <i>RIP</i> , <i>SNMP</i> , <i>TFTP</i> , <i>TIME</i> , and <i>WHO</i> .
<i>ICMP Type</i>	IPv6 ACL rule to match on the specified ICMP message type. This option is available only if the protocol is ICMPv6.
<i>ICMP Code</i>	IPv6 ACL rule to match on the specified ICMP message code. This option is available only if the protocol is ICMPv6.
<i>ICMP Message</i>	IPv6 ACL rule to match on the ICMP message type and code. Specify one of the following supported ICMPv6 messages: <i>Destination-Unreachable</i> , <i>Echo-Request</i> , <i>Echo-Reply</i> , <i>Header</i> , <i>Hop-Limit</i> , <i>MLD-Query</i> , <i>MLD-Reduction</i> , <i>MLD-Report</i> , <i>ND-NA</i> , <i>ND-NS</i> , <i>Next-Header</i> , <i>No-Admin</i> , <i>No-Route</i> , <i>Packet-Too-Big</i> , <i>Port-Unreachable</i> , <i>Router-Solicitation</i> , <i>Router-Advertisement</i> , <i>Router-Renumbering</i> , <i>Time-Exceeded</i> , and <i>Unreachable</i> . This option is available only if the protocol is ICMPv6.
<i>TCP Flags</i>	IPv6 ACL rule to match on the TCP flags. When a + flag is specified, a match occurs if the flag is set in the TCP header. When a - flag is specified, a match occurs if the flag is not set in the TCP header. When <i>Established</i> is specified, a match occurs if either RST or ACK bits are set in the TCP header. This option is available only if the protocol is TCP.
<i>Flow Label</i>	A 20-bit number that is unique to an IPv6 packet, used by end stations to signify quality-of-service handling in routers.

Access Control List Configuration Fields (Continued)

Field	Description
<i>IP DSCP</i>	The IP DSCP value in the IPv6 packet to match to the rule. The DSCP value is defined as the high-order six bits of the Service Type octet in the IPv6 header.
<i>Routing</i>	IPv6 ACL rule to match on routed packets.
<i>Match Criteria (MAC ACLs)</i> – The fields in this section specify the criteria to use to determine whether an Ethernet frame matches the rule. The fields described below apply to MAC ACLs.	
<i>Every</i>	When this option is selected, all packets will match the rule and will be either permitted or denied. This option is exclusive to all other match criteria – if <i>Every</i> is selected, no other match criteria can be configured. To configure specific match criteria, this option must be cleared.
<i>CoS</i>	The 802.1p user priority value to match within the Ethernet frame.
<i>Ethertype</i>	The EtherType value to match in an Ethernet frame. Specify the number associated with the EtherType or specify one of the following keywords: <i>AppleTalk</i> , <i>ARP</i> , <i>IBM SNA</i> , <i>IPv4</i> , <i>IPv6</i> , <i>IPX</i> , <i>MPLS</i> , <i>Unicast</i> , <i>NETBIOS</i> , <i>NOVELL</i> , <i>PPPoE</i> , or <i>RARP</i> .
<i>Source MAC Address / Mask</i>	The MAC address to match to an Ethernet frame's source port MAC address. If desired, enter the MAC mask associated with the source MAC to match. The MAC address mask specifies which bits in the source MAC to compare against an Ethernet frame, and uses F's and 0's in a wildcard format. An F means that the bit is not checked, and a 0 in a bit position means that the data must equal the value given for that bit. For example, if the MAC address is <i>aa:bb:cc:dd:ee:ff</i> , and the mask is <i>00:00:ff:ff:ff:ff</i> , all MAC addresses with <i>aa:bb:xx:xx:xx:xx</i> result in a match (where x is any hexadecimal number).
<i>Destination MAC Address / Mask</i>	The MAC address to match to an Ethernet frame's destination port MAC address. If desired, enter the MAC Mask associated with the destination MAC to match. The MAC address mask specifies which bits in the destination MAC to compare against an Ethernet frame. Use F's and 0's in the MAC mask, which is in a wildcard format. An F means that the bit is not checked, and a 0 in a bit position means that the data must equal the value given for that bit. For example, if the MAC address is <i>aa:bb:cc:dd:ee:ff</i> , and the mask is <i>00:00:ff:ff:ff:ff</i> , all MAC addresses with <i>aa:bb:xx:xx:xx:xx</i> result in a match (where x is any hexadecimal number).
<i>VLAN</i>	The VLAN ID to match within the Ethernet frame.
<i>Rule Attributes</i> – The fields in this section provide information about the actions to take on a frame or packet that matches the rule criteria. The attributes specify actions other than the basic Permit or Deny actions.	
<i>Assign Queue</i>	The number that identifies the hardware egress queue that will handle all packets matching this rule.
<i>Interface</i>	The interface to use for the action: Redirect Allows traffic that matches a rule to be redirected to the selected interface instead of being processed on the original port. The redirect function and mirror function are mutually exclusive. Mirror Allows traffic that matches a rule to be mirrored to a selected interface. Mirroring is similar to the redirect function, except that in flow-based mirroring a copy of the permitted traffic is delivered to the mirror interface while the packet itself is forwarded normally through the device.
<i>Log</i>	When this option is selected, logging is enabled for this ACL rule (subject to resource availability in the device). If the Access List Trap Flag is also enabled, this will cause periodic traps to be generated indicating the number of times this rule went into effect during the current report interval. A fixed 5 minute report interval is used for the entire system. A trap is not issued if the ACL rule hit count is zero for the current interval.
<i>Time Range Name</i>	The name of the time range that will impose a time limitation on the ACL rule. If a time range with the specified name does not exist, and the ACL containing this ACL rule is associated with an interface, the ACL rule is applied immediately. If a time range with specified name exists, and the ACL containing this ACL rule is associated with an interface, the ACL rule is applied when the specified time-range becomes active. The ACL rule is removed when the specified time-range with becomes inactive.
<i>Committed Rate / Burst Size</i>	The allowed transmission rate for frames on the interface (Committed Rate), and the number of bytes allowed in a temporary traffic burst (Burst Rate).

Use the buttons to perform the following tasks:


- To add an ACL rule entry, select the ID of the ACL that will include the rule from the *ACL Identifier* drop-down menu. Then, click **Add Rule** and configure the rule criteria and attributes (new rules cannot be created if the maximum number of rules has been reached). Finally, click **Submit** to apply the changes.
- To remove the most recently configured rule for an ACL, select the ID of the appropriate ACL from the ACL Identifier menu and click **Remove Last Rule**. You must confirm the action before the entry is deleted.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Access Control List Interface Summary

Use the *Access Control List Interface Summary* page to associate one or more ACLs with one or more interfaces on the device. When an ACL is associated with an interface, traffic on the port is checked against the rules defined within the ACL until a match is found. If the traffic does not match any rules within an ACL, it is dropped because of the implicit deny all rule at the end of each ACL.

To display the page, click **QoS > Access Control Lists > Interfaces** in the navigation menu.



Access Control List Interface Summary

Access Control List Interface Summary Fields

Field	Description
<i>Interface</i>	The interface that has an associated ACL.
<i>Direction</i>	Indicates whether the packet is checked against the rules in an ACL when it is received on an interface (<i>Inbound</i>) or after it has been received, routed, and is ready to exit an interface (<i>Outbound</i>).
<i>Sequence Number</i>	The order the ACL is applied to traffic on the interface relative to other ACLs associated with the interface in the same direction. When multiple ACLs are applied to the same interface in the same direction, the ACL with the lowest sequence number is applied first, and the other ACLs are applied in ascending numerical order.
<i>ACL Type</i>	The ACL type, which determines what criteria can be used to match packets. The type also determines which attributes can be applied to matching traffic. IPv4 ACLs classify Layer-3 and Layer-4 IPv4 traffic, IPv6 ACLs classify Layer-3 and Layer-4 IPv6 traffic, and MAC ACLs classify Layer-2 traffic. The ACL types are as follows: IPv4 Standard Match criteria is based on the source address of IPv4 packets. IPv4 Extended Match criteria can be based on the source and destination addresses, source and destination Layer-4 ports, and protocol type of IPv4 packets. IPv4 Named Match criteria is the same as IPv4 Extended ACLs, but the ACL ID can be an alphanumeric name instead of a number. IPv6 Named Match criteria can be based on information including the source and destination IPv6 addresses, source and destination Layer-4 ports, and protocol type within IPv6 packets. Extended MAC Match criteria can be based on the source and destination MAC addresses, 802.1p user priority, VLAN ID, and EtherType value within Ethernet frames.
<i>ACL Identifier</i>	The name or number that identifies the ACL. When applying an ACL to an interface, the <i>ACL Identifier</i> menu includes only the ACLs within the selected <i>ACL Type</i> .

Use the buttons to perform the following tasks:

- To apply an ACL to an interface, click **Add** and configure the settings in the available fields.
- To remove the association between an interface and an ACL, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Access Control List VLAN Summary

Use this page to associate one or more ACLs with one or more VLANs on the device.

To display the *Access Control List VLAN Summary* page, click **QoS > Access Control Lists > Interfaces** in the navigation menu.



Access Control List VLAN Summary

Access Control List VLAN Summary Fields

Field	Description
<i>VLAN ID</i>	The ID of the VLAN associated with the rest of the data in the row. When associating a VLAN with an ACL, use this field to select the desired VLAN.
<i>Direction</i>	Indicates whether the packet is checked against the rules in an ACL when it is received on a VLAN (<i>Inbound</i>) or after it has been received, routed, and is ready to exit a VLAN (<i>Outbound</i>).
<i>Sequence Number</i>	The order the ACL is applied to traffic on the VLAN relative to other ACLs associated with the VLAN in the same direction. When multiple ACLs are applied to the same VLAN in the same direction, the ACL with the lowest sequence number is applied first, and the other ACLs are applied in ascending numerical order.
<i>ACL Type</i>	The ACL type, which determines what criteria can be used to match packets. The type also determines which attributes can be applied to matching traffic. IPv4 ACLs classify Layer-3 and Layer-4 IPv4 traffic, IPv6 ACLs classify Layer-3 and Layer-4 IPv6 traffic, and MAC ACLs classify Layer-2 traffic. The ACL types are as follows: IPv4 Standard Match criteria is based on the source address of IPv4 packets. IPv4 Extended Match criteria can be based on the source and destination addresses, source and destination Layer-4 ports, and protocol type of IPv4 packets. IPv4 Named Match criteria is the same as IPv4 Extended ACLs, but the ACL ID can be an alphanumeric name instead of a number. IPv6 Named Match criteria can be based on information including the source and destination IPv6 addresses, source and destination Layer-4 ports, and protocol type within IPv6 packets. Extended MAC Match criteria can be based on the source and destination MAC addresses, 802.1p user priority, VLAN ID, and EtherType value within Ethernet frames.
<i>ACL Identifier</i>	The name or number that identifies the ACL. The permitted identifier depends on the ACL type. Standard and Extended IPv4 ACLs use numbers within a set range, and Named IPv4, IPv6, and MAC ACLs use alphanumeric characters.

Use the buttons to perform the following tasks:

- To associate an ACL with a VLAN, click **Add**, configure the settings, and click **Submit** to apply the changes.
- To remove the association between a VLAN and an ACL, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Configuring Auto VoIP

Voice over Internet Protocol (VoIP) allows you to make telephone calls using a computer network over a data network like the Internet. With the increased prominence of delay-sensitive applications (voice, video, and other multimedia applications) deployed in networks today, proper QoS configuration will ensure high-quality application performance. The Auto VoIP feature is intended to provide an easy classification mechanism for voice packets so that they can be prioritized above data packets in order to provide better QoS.

The Auto-VoIP feature explicitly matches VoIP streams in Ethernet switches and provides them with a better class of service than ordinary traffic. If you enable the Auto-VoIP feature on an interface, the interface scans incoming traffic for the following call-control protocols:

- Session Initiation Protocol (SIP)
- H.323
- Skinny Client Control Protocol (SCCP)

When a call-control protocol is detected the switch assigns the traffic in that session to the highest CoS queue, which is generally used for time-sensitive traffic.

Auto VoIP Global Configuration

Use the *Auto VoIP Global Configuration* page to configure the VLAN ID for the Auto VoIP VLAN or to reset the current Auto VoIP VLAN ID to the default value. Voice over Internet Protocol (VoIP) enables telephone calls over a data network. Because voice traffic is typically more time-sensitive than data traffic, the Auto VoIP feature helps provide a classification mechanism for voice packets so that they can be prioritized above data packets in order to provide better Quality of Service (QoS). With the Auto VoIP feature, voice prioritization is provided based on call-control protocols (SIP, SCCP, H.323) and/or OUI bits. When the device identifies voice traffic, it is placed in the VLAN specified on this page. The Auto VoIP feature does not rely on LLDP-MED support in connected devices.

To display this page, click **QoS > Auto VoIP > Global** in the navigation menu.

Auto VoIP Global Configuration

Auto VoIP Global Configuration Fields

Field	Description
<i>Auto VoIP VLAN</i>	The VLAN used to segregate VoIP traffic from other non-voice traffic.

Use the buttons to perform the following tasks:


- If you change the *Auto VoIP VLAN* field, click **Submit** to apply the change.
- To reset the VLAN to the default Auto VoIP VLAN, click **Reset**, and confirm the action.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

OUI Table Summary

Use the *OUI Table Summary* page to add and remove Organizationally Unique Identifiers (OUIs) from the OUI database the device maintains. Device hardware manufacturers can include an OUI in a network adapter to help identify the device. The OUI is a unique 24-bit number assigned by the IEEE registration authority. Several default OUIs have been preconfigured in the OUI database on the device.

To display the page, click **Quality of Service > Auto VoIP > OUI Table** in the navigation menu.



OUI Table Summary

OUI Table Summary Fields

Field	Description
<i>Telephony OUI</i>	The unique OUI that identifies the device manufacturer or vendor. The OUI is specified in three octet values (each octet is represented as two hexadecimal digits) separated by colons.
<i>Status</i>	Identifies whether the OUI is preconfigured on the system (<i>Default</i>) or added by a user (<i>Configured</i>).
<i>Description</i>	Identifies the manufacturer or vendor associated with the OUI.

Use the buttons to perform the following tasks:

- To add an OUI, click **Add**, specify an OUI and its description in the available fields, and click **Submit** to apply the changes.
- To remove one or more configured OUIs, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

OUI Based Auto VoIP

Use this page to configure the Organizationally Unique Identifier (OUI) based Auto VoIP priority and to enable or disable the Auto VoIP mode on the interfaces.

To display the *OUI Based Auto VoIP* page, click **QoS** > **Auto VoIP** > **OUI Based Auto VoIP** in the navigation menu.

The screenshot shows the 'OUI Based Auto VoIP' configuration page. At the top, there are tabs: 'Global', 'OUI Table', 'OUI Based Auto VoIP' (selected), and 'Protocol Based Auto VoIP'. Below the tabs, there are two input fields: 'Auto VoIP VLAN' with the value 'Not Configured' and 'Priority' with the value '7' (range 0 to 7). A 'Display' dropdown is set to '10 rows', and it says 'Showing 1 to 10 of 32 entries'. A 'Filter' input field is on the right. The main table has columns: 'Interface', 'Auto VoIP Mode', and 'Operational Status'. The table lists 10 interfaces from q/1 to q/10, all with 'Disable' mode and 'Down' status. At the bottom, there are buttons: 'First', 'Previous', '1', '2', '3', '4', 'Next', 'Last', 'Submit', 'Refresh', 'Edit All', and 'Cancel'.

OUI Based Auto VoIP

OUI Based Auto VoIP Fields

Field	Description
<i>Auto VoIP VLAN</i>	The VLAN used to segregate VoIP traffic from other non-voice traffic. All VoIP traffic that matches a value in the known OUI list gets assigned to this VoIP VLAN.
<i>Priority</i>	The 802.1p priority used for traffic that matches a value in the known OUI list. If the Auto VoIP mode is enabled and the interface detects an OUI match, the device assigns the traffic in that session to the traffic class mapped to this priority value. Traffic classes with a higher value are generally used for time-sensitive traffic.
<i>Interface</i>	The interface associated with the rest of the data in the row. When editing Auto VoIP settings on one or more interfaces, this field identifies the interface(s) being configured.
<i>Auto VoIP Mode</i>	The administrative mode of OUI-based Auto VoIP on the interface.
<i>Operational Status</i>	The interface's operational status (<i>Up</i> or <i>Down</i>). To be <i>Up</i> , an interface must be administratively enabled and have a link.

Use the buttons to perform the following tasks:

- If you change the *Priority* field, click **Submit** to apply the change.
- To configure settings on one or more interfaces, select each interface and click **Edit**. In the *Edit OUI Based Port Configuration* window, edit the settings as needed, and click **Submit** to apply the changes.

- To configure settings on all interfaces, click **Edit All**. In the *Edit OUI Based Port Configuration* window, change the settings as needed, and click **Submit** to apply the changes.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Protocol Based Auto VoIP

Use this page to configure the protocol-based Auto VoIP priority settings and to enable or disable the protocol-based Auto VoIP mode on the interfaces.

To display the *Protocol Based Auto VoIP* page, click **QoS > Auto VoIP > Protocol Based Auto VoIP** in the navigation menu. A portion of the UI page is shown below.

The screenshot shows the 'Protocol Based Auto VoIP' configuration page. At the top, there are tabs for 'Global', 'OUI Table', 'OUI Based Auto VoIP', and 'Protocol Based Auto VoIP'. The 'Protocol Based Auto VoIP' tab is selected. Below the tabs, there are configuration fields: 'Auto VoIP VLAN' (Not Configured), 'Prioritization Type' (Remark and Traffic Class radio buttons, with Traffic Class selected), '802.1p Priority' (a text box with '00 to 71'), and 'Traffic Class' (a text box with '7' and '00 to 71'). Below these fields, there is a 'Display' dropdown set to '10 rows' and a 'Showing 1 to 10 of 32 entries' indicator. A table with 3 columns: 'Interface', 'Auto VoIP Mode', and 'Operational Status' is shown. The table has 11 rows, each with a checkbox in the 'Interface' column, the interface name (0/1 to 0/10), 'Disable' in the 'Auto VoIP Mode' column, and 'Down' in the 'Operational Status' column. At the bottom of the table, there are buttons for 'First', 'Previous', '1', '2', '3', '4', 'Next', and 'Last'. Below the table, there are buttons for 'Submit', 'Refresh', 'Edit All', and 'Cancel'. At the very bottom, there is a copyright notice: '© Copyright 2013-2014 Ubiquiti Networks, Inc.'

Protocol Based Auto VoIP

Protocol Based Auto VoIP Fields

Field	Description
<i>Auto VoIP VLAN</i>	The VLAN used to segregate VoIP traffic from other non-voice traffic. All VoIP traffic in a session identified by the call-control protocol gets assigned to this VoIP VLAN.
<i>Prioritization Type</i>	The method used to prioritize VoIP traffic when a call-control protocol is detected, which is one of the following: • Remark Remark the voice traffic with the specified 802.1p priority value at the ingress interface. • Traffic Class Assign VoIP traffic to the specified traffic class when egressing the interface.

Protocol Based Auto VoIP Fields (Continued)

Field	Description
802.1p Priority	The 802.1p priority used for protocol-based VoIP traffic. This field can be configured if the <i>Prioritization Type</i> is 802.1p Priority. If the <i>Auto VoIP Mode</i> is enabled and the interface detects a call-control protocol, the device marks traffic in that session with the specified 802.1p priority value to ensure voice traffic always gets the highest priority throughout the network path. Egress tagging must be administratively enabled on the appropriate uplink port to carry the remarked priority at the egress port.
Traffic Class	The traffic class used for protocol-based VoIP traffic. This field can be configured if the <i>Prioritization Type</i> is Traffic Class . If the <i>Auto VoIP Mode</i> is enabled and the interface detects a call-control protocol, the device assigns the traffic in that session to the configured Class of Service (CoS) queue. Traffic classes with a higher value are generally used for time-sensitive traffic. The CoS queue associated with the specified traffic class should be configured with the appropriate bandwidth allocation to allow priority treatment for VoIP traffic.
Interface	The interface associated with the rest of the data in the row. When editing Auto VoIP settings on one or more interfaces, this field identifies the interface(s) being configured.
Auto VoIP Mode	The administrative mode of the Auto VoIP feature on the interface: • Enable The interface scans incoming traffic for the following call-control protocols: • Session Initiation Protocol (SIP) • H.323 • Skinny Client Control Protocol (SCCP) • Disable The interface does not use the Auto VoIP feature to scan for call-control protocols.
Operational Status	The operational status of an interface. To be up, an interface must be administratively enabled and have a link.

Use the buttons to perform the following tasks:

- If you edit any fields, click **Submit** to apply the changes.
- To configure settings on one or more interfaces, select each interface and click **Edit**. In the *Edit Protocol Based Port Configuration* window, edit the settings as needed, and click **Submit** to apply the changes.
- To configure settings on all interfaces, click **Edit All**. In the *Edit Protocol Based Port Configuration* window, change the settings as needed, and click **Submit** to apply the changes.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Configuring Class of Service

The Class of Service (CoS) queueing feature lets you directly configure certain aspects of switch queueing. This provides the desired QoS behavior for different types of network traffic when the complexities of DiffServ are not required. The priority of a packet arriving at an interface can be used to steer the packet to the appropriate outbound CoS queue through a mapping table. CoS queue characteristics that affect queue mapping, such as minimum guaranteed bandwidth, transmission rate shaping, etc., are user-configurable at the queue (or port) level.

Seven queues per port are supported. Although the hardware supports eight queues, one queue is always reserved for internal use by the stacking subsystem.

CoS IP DSCP Mapping Configuration

Use the *CoS IP DSCP Mapping Configuration* page to map an IP DSCP value to an internal traffic class.

To display the page, click **QoS > Class of Service > IP DSCP** in the navigation menu.

IP DSCP	Traffic Class
0	0
1	1
2	1
3	1
4	1
5	1
6	1
7	1
8	0
9	0
10	0
11	0
12	0
13	0
14	0
15	0
16	7
17	7
18	7
19	7
20	7
21	7
22	7
23	7

CoS IP DSCP Mapping Configuration – 1 of 2

24	☐ 0	☒ 1	☐ 2	☐ 3	☐ 4	☐ 5	☐ 6	☐ 7
25	☐ 0	☒ 1	☐ 2	☐ 3	☐ 4	☐ 5	☐ 6	☐ 7
26	☐ 0	☒ 1	☐ 2	☐ 3	☐ 4	☐ 5	☐ 6	☐ 7
27	☐ 0	☒ 1	☐ 2	☐ 3	☐ 4	☐ 5	☐ 6	☐ 7
28	☐ 0	☒ 1	☐ 2	☐ 3	☐ 4	☐ 5	☐ 6	☐ 7
29	☐ 0	☒ 1	☐ 2	☐ 3	☐ 4	☐ 5	☐ 6	☐ 7
30	☐ 0	☒ 1	☐ 2	☐ 3	☐ 4	☐ 5	☐ 6	☐ 7
31	☐ 0	☒ 1	☐ 2	☐ 3	☐ 4	☐ 5	☐ 6	☐ 7
32	☐ 0	☐ 1	☒ 2	☐ 3	☐ 4	☐ 5	☐ 6	☐ 7
33	☐ 0	☐ 1	☒ 2	☐ 3	☐ 4	☐ 5	☐ 6	☐ 7
34	☐ 0	☐ 1	☒ 2	☐ 3	☐ 4	☐ 5	☐ 6	☐ 7
35	☐ 0	☐ 1	☒ 2	☐ 3	☐ 4	☐ 5	☐ 6	☐ 7
36	☐ 0	☐ 1	☒ 2	☐ 3	☐ 4	☐ 5	☐ 6	☐ 7
37	☐ 0	☐ 1	☒ 2	☐ 3	☐ 4	☐ 5	☐ 6	☐ 7
38	☐ 0	☐ 1	☒ 2	☐ 3	☐ 4	☐ 5	☐ 6	☐ 7
39	☐ 0	☐ 1	☒ 2	☐ 3	☐ 4	☐ 5	☐ 6	☐ 7
40	☐ 0	☐ 1	☒ 2	☐ 3	☐ 4	☐ 5	☐ 6	☐ 7
41	☐ 0	☐ 1	☒ 2	☐ 3	☐ 4	☐ 5	☐ 6	☐ 7
42	☐ 0	☐ 1	☒ 2	☐ 3	☐ 4	☐ 5	☐ 6	☐ 7
43	☐ 0	☐ 1	☒ 2	☐ 3	☐ 4	☐ 5	☐ 6	☐ 7
44	☐ 0	☐ 1	☒ 2	☐ 3	☐ 4	☐ 5	☐ 6	☐ 7
45	☐ 0	☐ 1	☒ 2	☐ 3	☐ 4	☐ 5	☐ 6	☐ 7
46	☐ 0	☐ 1	☒ 2	☐ 3	☐ 4	☐ 5	☐ 6	☐ 7
47	☐ 0	☐ 1	☒ 2	☐ 3	☐ 4	☐ 5	☐ 6	☐ 7
48	☐ 0	☐ 1	☐ 2	☒ 3	☐ 4	☐ 5	☐ 6	☐ 7
49	☐ 0	☐ 1	☐ 2	☒ 3	☐ 4	☐ 5	☐ 6	☐ 7
50	☐ 0	☐ 1	☐ 2	☒ 3	☐ 4	☐ 5	☐ 6	☐ 7
51	☐ 0	☐ 1	☐ 2	☒ 3	☐ 4	☐ 5	☐ 6	☐ 7
52	☐ 0	☐ 1	☐ 2	☒ 3	☐ 4	☐ 5	☐ 6	☐ 7
53	☐ 0	☐ 1	☐ 2	☒ 3	☐ 4	☐ 5	☐ 6	☐ 7
54	☐ 0	☐ 1	☐ 2	☒ 3	☐ 4	☐ 5	☐ 6	☐ 7
55	☐ 0	☐ 1	☐ 2	☒ 3	☐ 4	☐ 5	☐ 6	☐ 7
56	☐ 0	☐ 1	☐ 2	☒ 3	☐ 4	☐ 5	☐ 6	☐ 7
57	☐ 0	☐ 1	☐ 2	☒ 3	☐ 4	☐ 5	☐ 6	☐ 7
58	☐ 0	☐ 1	☐ 2	☒ 3	☐ 4	☐ 5	☐ 6	☐ 7
59	☐ 0	☐ 1	☐ 2	☒ 3	☐ 4	☐ 5	☐ 6	☐ 7
60	☐ 0	☐ 1	☐ 2	☒ 3	☐ 4	☐ 5	☐ 6	☐ 7
61	☐ 0	☐ 1	☐ 2	☒ 3	☐ 4	☐ 5	☐ 6	☐ 7
62	☐ 0	☐ 1	☐ 2	☒ 3	☐ 4	☐ 5	☐ 6	☐ 7
63	☐ 0	☐ 1	☐ 2	☒ 3	☐ 4	☐ 5	☐ 6	☐ 7

© Copyright 2013-2014 Ubiquiti Networks, Inc.

CoS IP DSCP Mapping Configuration – 2 of 2

CoS IP DSCP Mapping Configuration Fields

Field	Description
Interface	The interface to configure. To configure the same IP DSCP-to-Traffic Class mappings on all interfaces, select Global .
IP DSCP	The list of possible IP DSCP values the IP header can include.
Traffic Class	The internal traffic class to which the corresponding IP DSCP priority value is mapped. The higher the traffic class value, the higher its priority is for sending traffic.

Use the buttons to perform the following tasks:

- If you change any fields, click **Submit** to apply the changes.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

Interface Configuration

Use the *CoS Interface Configuration* page to apply an interface shaping rate to all ports or to a specific port.

To display the *CoS Interface Configuration* page, click **QoS** > **Class of Service** > **Interface** in the navigation menu.

CoS Interface Configuration

CoS Interface Configuration Fields

Field	Description
Interface	The interface to configure. To configure the same settings on all interfaces, select Global .
Trust Mode	The trust mode for ingress traffic on the interface, which is one of the following: • untrusted The interface ignores any priority designations encoded in incoming packets, and instead sends the packets to a traffic queue based on the ingress port's default priority. • trust dot1p The port accepts at face value the 802.1p priority designation encoded within packets arriving on the port. • trust ip-dscp The port accepts at face value the IP DSCP priority designation encoded within packets arriving on the port.
Interface Shaping Rate	The upper limit on how much traffic can leave a port. The limit on maximum transmission bandwidth has the effect of smoothing temporary traffic bursts over time so that the transmitted traffic rate is bounded. The specified value represents a percentage of the maximum negotiated bandwidth.
WRED Decay Exponent	The decay exponent value used with the Weighted Random Early Detection (WRED) average queue length calculation algorithm.

Use the buttons to perform the following tasks:

- If you change any fields on the page, click **Submit** to apply the changes.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

CoS Interface Queue Configuration

Use the *CoS Interface Queue Configuration* page to define what a particular queue does by configuring switch egress queues. User-configurable parameters control the amount of bandwidth used by the queue, the queue depth during times of congestion, and the scheduling of packet transmission from the set of all queues on a port. Each port has its own CoS queue-related configuration. The configuration process is simplified by allowing each CoS queue parameter to be configured globally or per-port. A global configuration change is automatically applied to all ports in the system.

To display the page, click **QoS > Class of Service > Queue** in the navigation menu.

The screenshot shows the 'CoS Interface Queue Configuration' page. At the top, there are four tabs: 'IP DSCP', 'Interface', 'Queue', and 'Drop Precedence'. The 'Queue' tab is selected. Below the tabs, there's a section for 'Interface' with a dropdown menu showing 'Q/L'. Below that, 'Total Minimum Bandwidth Allocation (G)' is set to '0'. The main part of the page is a table with columns: 'Queue ID', 'Minimum Bandwidth (G)', 'Scheduler Type', and 'Queue Management Type'. There are 8 rows, numbered 0 to 7. Queue 6 is highlighted in blue and has a checkmark in the first column. All 'Minimum Bandwidth' values are '0', and all 'Scheduler Type' values are 'Weighted'. All 'Queue Management Type' values are 'TailDrop'. At the bottom, there are three buttons: 'Refresh', 'Edit', and 'Restore Default'. A copyright notice '© Copyright 2013-2014 Ubiquiti Networks, Inc.' is visible in the bottom right corner of the interface.

CoS Interface Queue Configuration

CoS Interface Queue Configuration Fields

Field	Description
<i>Interface</i>	Specifies the interface (physical, LAG, or Global) to configure.
<i>Total Minimum Bandwidth Allocation (%)</i>	Shows the total minimum bandwidth allocated to the selected interface for all the queues.
<i>Queue ID</i>	The CoS queue. The higher the queue value, the higher its priority is for sending traffic.
<i>Minimum Bandwidth</i>	The minimum guaranteed bandwidth allocated to the selected queue on the interface. Setting this value higher than its corresponding <i>Maximum Bandwidth</i> automatically increases the maximum to the same value. A value of 0 (zero) means no guaranteed minimum. The sum of individual <i>Minimum Bandwidth</i> values for all queues in the selected interface cannot exceed defined a maximum of 100.
<i>Scheduler Type</i>	The type of queue processing. Defining this value on a per-queue basis allows you to create the desired service characteristics for different types of traffic. The options are as follows: Weighted Weighted round robin associates a weight to each queue. This is the default. Strict Strict priority services traffic with the highest priority on a queue first.
<i>Queue Management Type</i>	The type of queue depth management techniques used for all queues on this interface. Options are: Taildrop All packets on a queue are safe until congestion occurs. At this point, any additional packets queued are dropped. WRED Weighted Random Early Detection (WRED) drops packets selectively based on their drop precedence level.

CoS Interface Queue Drop Precedence Configuration

Use this page to configure the queue drop precedence on a per-queue, per-interface basis. When an interface is configured with taildrop queue management, all packets on a queue are safe until congestion occurs. If congestion occurs, any additional packets queued are dropped. Weighted Random Early Detection (WRED) drops packets selectively based their drop precedence level.

To display the *CoS Interface Queue Drop Precedence Configuration* page, click **QoS > Class of Service > Drop Precedence** in the navigation menu.

CoS Interface Queue Drop Precedence Configuration

CoS Interface Queue Drop Precedence Configuration Fields

Field	Description
<i>Interface</i>	The interface on which to configure the queue drop precedence settings. To configure the same settings on all interfaces, select the <i>Global</i> menu option.
<i>Queue ID</i>	The CoS queue on which to configure the drop precedence settings. The higher the queue value, the higher its priority is for sending traffic.
<i>Drop Precedence Level</i>	The four drop precedence levels.
<i>WRED Minimum Threshold</i>	The minimum queue threshold below which now packets are dropped for the associated drop precedence level. After the minimum is reached, WRED randomly drops packets based on their priority (DSCP or IP precedence). This setting applies to the interface if it is configured with a WRED queue management type.
<i>WRED Maximum Threshold</i>	The maximum queue threshold above which all packets are dropped for the associated drop precedence level. After the maximum is reached, WRED drops all packets based on their priority (DSCP or IP precedence). This setting applies to the interface if it is configured with a WRED queue management type.
<i>WRED Drop Probability Scale</i>	The packet drop probability for the drop precedence level. This setting applies to the interface if it is configured with a WRED queue management type.

Use the buttons to perform the following tasks:

- To edit the settings on an interface, select the entry from table, and click **Edit**. When you have completed the changes, click **Submit** to apply the changes.
- Click **Restore Defaults** to restore all drop precedence settings on the selected interface to the default values. If **Global** is selected in the *Interface* field, all default settings for all interfaces are restored.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Configuring Diffserv

Use this page to configure the administrative mode of Differentiated Services (DiffServ) support on the device and to view the current and maximum number of entries in each of the main DiffServ private MIB tables. DiffServ allows traffic to be classified into streams and given certain QoS treatment in accordance with defined per-hop behaviors.

Packets are classified and processed based on defined criteria. The classification criteria is defined by a class. The processing is defined by a policy's attributes. Policy attributes may be defined on a per-class instance basis, and it is these attributes that are applied when a match occurs. A policy can contain multiples classes. When the policy is active, the actions taken depend on which class matches the packet.

Diffserv Global Configuration and Status

Use the *Diffserv Global Configuration and Status* page to configure the Global DiffServ settings on the device.

To display the page, click **QoS > Diffserv > Global** in the navigation menu.

Diffserv Global Configuration and Status

Diffserv Global Configuration and Status Fields

Field	Description
<i>Diffserv Admin Mode</i>	The administrative mode of DiffServ on the device. While disabled, the DiffServ configuration is retained and can be changed, but it is not active. While enabled, Differentiated Services are active.
<i>MIB Table</i> – The information in this table displays the number of entries (rows) that are currently in each of the main DiffServ private MIB tables and the maximum number of rows that can exist in each table.	
<i>Class Table</i>	The current and maximum number of classifier entries in the table. DiffServ classifiers differentiate among traffic types.
<i>Class Rule Table</i>	The current and maximum number of class rule entries in the table. Class rules specify the match criteria that belong to a class definition.
<i>Policy Table</i>	The current and maximum number of policy entries in the table. The policy determines the traffic conditioning or service provisioning actions applied to a traffic class.
<i>Policy Instance Table</i>	The current and maximum number of policy-class instance entries in the table. A policy-class instance is a policy that is associated with an existing DiffServ class.
<i>Policy Attribute Table</i>	The current and maximum number of policy attribute entries in the table. A policy attribute entry attaches various policy attributes to a policy-class instance.
<i>Service Table</i>	The current and maximum number of service entries in the table. A service entry associates a DiffServ policy with an interface and inbound or outbound direction.

Use the buttons to perform the following tasks:

- If you change any fields on this page, click **Submit** to apply the changes.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Diffserv Class Summary

Use this page to create or remove DiffServ classes and to view summary information about the classes that exist on the device. Creating a class is the first step in using DiffServ to provide Quality of Service. After a class is created, you can define the match criteria for the class.

To display the *Diffserv Class Summary* page, click **QoS > Diffserv > Class Summary** in the navigation menu.



Diffserv Class Summary

Diffserv Class Summary Fields

Field	Description
<i>Name</i>	The name of the DiffServ class. When adding a new class or renaming an existing class, the name of the class is specified in the <i>Class</i> field of the dialog window.
<i>Type</i>	The class type, which is one of the following: All All the various match criteria defined for the class should be satisfied for a packet match. All signifies the logical AND of all the match criteria.
<i>Protocol</i>	The Layer-3 protocol to use for filtering class types, which is either <i>IPv4</i> or <i>IPv6</i> .
<i>Match Criteria</i>	The criteria used to match packets.

Use the buttons to perform the following tasks:

- To add a DiffServ class, click **Add**.
- To change the name of an existing class, select the entry to modify and click **Rename**.
- To remove one or more configured classes, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.
- Click **Refresh** to update the page with the most current data from the switch.

Diffserv Class Configuration

Use this page to define the criteria to associate with a DiffServ class. As packets are received or transmitted, these DiffServ classes are used to classify and prioritize packets. Each class can contain multiple match criteria.

After you select the class to configure from the Class menu, use the buttons to perform the following tasks:

To define criteria for matching packets within a class, click **Add Match Criteria**. Once you add a match criteria entry to a class, you cannot edit or remove the entry. However, you can add more match criteria entries to a class until the maximum number of entries has been reached for the class. This button is not available if the class type is ACL because the match criteria are defined by the ACL rules.

To remove the associated reference class from the selected class, click **Remove Reference Class**. Note that unless the reference class is the last entry in the list of match criteria, the Reference Class match type remains in the list as a placeholder, but the associated value is N/A, and the previously referenced class is removed.

To display the *Diffserv Class Configuration* page, click **QoS > Diffserv > Class Configuration** in the navigation menu.

Diffserv Class Configuration

Diffserv Class Configuration Fields

Field	Description
<i>Class</i>	The name of the class. To configure match criteria for a class, select its name from the menu.
<i>Type</i>	The class type, which is one of the following: All All the various match criteria defined for the class should be satisfied for a packet match. All signifies the logical AND of all the match criteria.
<i>Protocol</i>	The Layer-3 protocol to use for filtering class types, which is either IPv4 or IPv6.
<i>Match Criteria</i>	The type of match criteria defined for the selected class. If the <i>Type</i> is <i>ACL</i> , no information about the match criteria is available on this page.
<i>Value</i>	The configured value of the match criteria that corresponds to the match type.
<i>Add Match Criteria</i> window – After you click Add Match Criteria, this window opens and allows you to define the match criteria for the selected class. The window lists the match criteria available for the class. To add match criteria, select the check box associated with the criteria type. The fields to configure the match values appear after you select the match type. Each match criteria type can be used only once within a class. If a reference class includes the match criteria type, it cannot be used as an additional match type within the class, and the match criteria type cannot be selected or configured.	
<i>Any</i>	Select this option to specify that all packets are considered to match the specified class. There is no need to configure additional match criteria if <i>Any</i> is selected because a match will occur on all packets.

Diffserv Class Configuration Fields (Continued)

Field	Description
<i>Reference Class</i>	Select this option to reference another class for criteria. The match criteria defined in the referenced class is as match criteria in addition to the match criteria you define for the selected class. After selecting this option, the classes that can be referenced are displayed. Select the class to reference. A class can reference at most one other class of the same type.
<i>Class of Service</i>	Select this option to require the Class of Service (CoS) value in an Ethernet frame header to match the specified CoS value.
<i>Secondary Class of Service</i>	Select this option to require the secondary CoS value in an Ethernet frame header to match the specified secondary CoS value.
<i>Ethertype</i>	Select this option to require the <i>EtherType</i> value in the Ethernet frame header to match the specified <i>EtherType</i> value. After you select this option, specify the <i>EtherType</i> value in one of these two fields: • <i>Ethertype Keyword</i> – The menu includes several common protocols that are mapped to their <i>EtherType</i> values. • <i>Ethertype Value</i> – This field accepts custom <i>EtherType</i> values.
<i>VLAN</i>	Select this option to require a packet's VLAN ID to match a VLAN ID or a VLAN ID within a continuous range. If you configure a range, a match occurs if a packet's VLAN ID is the same as any VLAN ID within the range. After you select this option, use the following fields to configure the VLAN match criteria: • VLAN ID Start The VLAN ID to match or the VLAN ID with the lowest value within a range of VLANs. • VLAN ID End The VLAN ID with the highest value within the range of VLANs. This field is not required if the match criteria is a single VLAN ID.
<i>Secondary VLAN</i>	Select this option to require a packet's VLAN ID to match a secondary VLAN ID or a secondary VLAN ID within a continuous range. If you configure a range, a match occurs if a packet's secondary VLAN ID is the same as any secondary VLAN ID within the range. After you select this option, use the following fields to configure the secondary VLAN match criteria: • Secondary VLAN ID Start The secondary VLAN ID to match or the secondary VLAN ID with the lowest value within a range of VLANs. • Secondary VLAN ID End The secondary VLAN ID with the highest value within the range of VLANs. This field is not required if the match criteria is a single VLAN ID.
<i>Source MAC Address</i>	Select this option to require a packet's source MAC address to match the specified MAC address. After you select this option, use the following fields to configure the source MAC address match criteria: • MAC Address The source MAC address to match. • MAC Mask The MAC mask, which specifies the bits in the source MAC address to compare against an Ethernet frame. Use F's and 0's to configure the MAC mask. An F means that the bit is checked, and a 0 in a bit position means that the data is not significant. For example, if the MAC address is <i>aa:bb:cc:dd:ee:ff</i>, and the mask is <i>ff:ff:00:00:00:00</i>, all MAC addresses with <i>aa:bb:xx:xx:xx:xx</i> result in a match (where x is any hexadecimal number). Note that this is not a wildcard mask, which ACLs use.
<i>Destination MAC Address</i>	Select this option to require a packet's destination MAC address to match the specified MAC address. After you select this option, use the following fields to configure the destination MAC address match criteria: • MAC Address The destination MAC address to match. • MAC Mask The MAC mask, which specifies the bits in the destination MAC address to compare against an Ethernet frame. Use F's and 0's to configure the MAC mask. An F means that the bit is checked, and a 0 in a bit position means that the data is not significant. For example, if the MAC address is <i>aa:bb:cc:dd:ee:ff</i>, and the mask is <i>ff:ff:00:00:00:00</i>, all MAC addresses with <i>aa:bb:xx:xx:xx:xx</i> result in a match (where x is any hexadecimal number). Note that this is not a wildcard mask, which ACLs use.
<i>Source IPv6 Address</i>	Select this option to require the source IPv6 address in a packet header to match the specified values. After you select this option, use the following fields to configure the source IPv6 address match criteria: • Source Prefix The source IPv6 prefix to match. • Source Prefix Length The IPv6 prefix length.
<i>Destination IPv6 Address</i>	Select this option to require the destination IPv6 address in a packet header to match the specified values. After you select this option, use the following fields to configure the destination IPv6 address match criteria: • Destination Prefix The destination IPv6 prefix to match. • Destination Prefix Length The IPv6 prefix length.

Diffserv Class Configuration Fields (Continued)

Field	Description
<i>Source L4 Port</i>	Select this option to require a packet's TCP/UDP source port to match the specified port or the port number within a range of port numbers. If you configure a range, a match occurs if a packet's source port number is the same as any source port number within the range. After you select this option, use the following fields to configure a source port keyword, source port number, or source port range for the match criteria: • Protocol Select the desired L4 keyword from the list on which the match is based. If you select a keyword, the other source port configuration fields are not available. • Port End A user-defined L4 source port number to match or the source port number with the lowest value within a range of ports. • Port Start The source port with the highest value within the range of ports. This field is not required if the match criteria is a single port.
<i>Destination L4 Port</i>	Select this option to require a packet's TCP/UDP destination port to match the specified port or the port number within a range of port numbers. If you configure a range, a match occurs if a packet's destination port number is the same as any destination port number within the range. After you select this option, use the following fields to configure a destination port keyword, destination port number, or destination port range for the match criteria: • Protocol Select the desired L4 keyword from the list on which the match is based. If you select a keyword, the other destination port configuration fields are not available. • Port End A user-defined L4 destination port number to match or the destination port number with the lowest value within a range of ports. • Port Start The destination port with the highest value within the range of ports. This field is not required if the match criteria is a single port.
<i>IP DSCP</i>	Select this option to require the packet's IP DiffServ Code Point (DSCP) value to match the specified value. The DSCP value is defined as the high-order six bits of the Service Type octet in the IP header. After you select this option, use one of the following fields to configure the IP DSCP match criteria: • IP DSCP Keyword The IP DSCP keyword code that corresponds to the IP DSCP value to match. If you select a keyword, you cannot configure an IP DSCP Value. • IP DSCP Value The IP DSCP value to match.
<i>IP Precedence</i>	Select this option to require the packet's IP Precedence value to match the number configured in the IP Precedence Value field. The IP Precedence field in a packet is defined as the high-order three bits of the Service Type octet in the IP header.
<i>IP TOS</i>	Select this option to require the packet's Type of Service (ToS) bits in the IP header to match the specified value. The IP ToS field in a packet is defined as all eight bits of the Service Type octet in the IP header. After you select this option, use the following fields to configure the ToS match criteria: • IP TOS Bits Enter a two-digit hexadecimal number to match the bits in a packet's ToS field. • IP TOS Mask Specify the bit positions used for comparison against the IP ToS field in a packet.
<i>Protocol</i>	Select this option to require a packet header's Layer-4 protocol to match the specified value. After you select this option, use one of the following fields to configure the protocol match criteria: • Protocol The L4 keyword that corresponds to value of the IANA protocol number to match. If you select a keyword, you cannot configure a Protocol Value. • Protocol Value The IANA L4 protocol number value to match.
<i>Flow Label</i>	Select this option to require an IPv6 packet's flow label to match the configured value. The flow label is a 20-bit number that is unique to an IPv6 packet, used by end stations to signify quality-of-service handling in routers.

Use the buttons to perform the following tasks:

- To define the match criteria for the selected class, click **Add Match Criteria**. In the *Add Match Criteria* window, configure the fields shown in the table below, and click **Submit** to apply the changes. Once you add a match criteria entry to a class, you cannot edit or remove the entry. However, you can add more match criteria entries to a class until the maximum number of entries has been reached for the class.
- To remove the associated reference class from the selected class, click **Remove Reference Class** and confirm the action.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Diffserv Policy Summary

Use this page to create or remove DiffServ policies and to view summary information about the policies that exist on the device. A policy defines the QoS attributes for one or more traffic classes. A policy attribute identifies the action taken when a packet matches a class rule. A policy is applied to a packet when a class match within that policy is found.

To display the *Diffserv Policy Summary* page, click **QoS > Diffserv > Policy Summary** in the navigation menu.



Diffserv Policy Summary

Diffserv Policy Summary Fields

Field	Description
<i>Name</i>	The name of the DiffServ policy. When adding a new policy or renaming an existing policy, the name of the policy is specified in the <i>Policy</i> field of the <i>Add Policy</i> dialog box.
<i>Type</i>	The traffic flow direction to which the policy is applied: In The policy is specific to inbound traffic. Out The policy is specific to outbound traffic direction.
<i>Member Classes</i>	The DiffServ class or classes that have been added to the policy.

Use the buttons to perform the following tasks:


- To add a DiffServ policy, click **Add**, configure the fields, and click **Submit** to apply the changes.
- To change the name of an existing policy, select the entry to modify and click **Rename**, enter the new name, and click **Submit** to apply the change.
- To remove one or more configured policies, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Diffserv Policy Configuration

Use the *Diffserv Policy Configuration* page to add or remove a DiffServ policy-class association and to configure the policy attributes. The policy attributes identify the action or actions taken when a packet matches a class rule.

To display the page, click **QoS > Diffserv > Policy Configuration** in the navigation menu.



Diffserv Policy Configuration

Diffserv Policy Configuration Fields

Field	Description
<i>Policy</i>	The name of the policy. To add a class to the policy, remove a class from the policy, or configure the policy attributes, you must first select its name from the menu.
<i>Type</i>	The traffic flow direction to which the policy is applied.
<i>Class</i>	The DiffServ class or classes associated with the policy. The policy is applied to a packet when a class match within that policy-class is found.
<i>Policy Attribute Details</i>	The policy attribute types and their associated values that are configured for the policy.
<i>Add Policy Attribute window</i> – Click Add Attribute to open this window and define the policy attributes for the selected policy. To add and configure policy attributes, select the check box for the attribute type and configure the fields for the attribute values.	
<i>Assign Queue</i>	Select this option to assign matching packets to a traffic queue. Use the Queue ID Value field to select the queue to which the packets of this policy-class are assigned.
<i>Drop</i>	Select this option to drop packets that match the policy-class.
<i>Mark CoS</i>	Select this option to mark all packets in a traffic stream with the specified Class of Service (CoS) queue value. Use the <i>Class of Service</i> field to select the CoS value to mark in the priority field of the 802.1p header (the only tag in a single tagged packet or the first or outer 802.1Q tag of a double VLAN tagged packet). If the packet does not already contain this header, one is inserted.
<i>Mark CoS as Secondary CoS</i>	Select this option to mark all packets in a traffic stream with the specified secondary CoS queue number. Use the <i>Class of Service</i> field to select the CoS value to mark in the priority field of the 802.1p header in the secondary (inner) 802.1Q tag of a double VLAN tagged packet. If the packet does not already contain this header, one is inserted.
<i>Mark IP DSCP</i>	Select this option to mark all packets in the associated traffic stream with the specified IP DSCP value. Then, use one of the following fields to configure the IP DSCP value to mark in packets that match the policy-class: IP DSCP Keyword The IP DSCP keyword code that corresponds to the IP DSCP value. If you select a keyword, you cannot configure an IP DSCP Value. IP DSCP Value The IP DSCP value.
<i>Mark IP Precedence</i>	Select this option to mark all packets in the associated traffic stream with the specified IP Precedence value. Then, select the <i>IP Precedence Value</i> to mark in packets that match the policy-class.
<i>Mirror Interface</i>	Select this option to copy the traffic stream to a specified egress port (physical or LAG) without bypassing normal packet forwarding. This can occur in addition to any marking or policing action. It may also be specified along with a QoS queue assignment. Use the Interface menu to select the interface to which traffic is mirrored.

Diffserv Policy Configuration Fields (Continued)

Field	Description
<i>Police Simple</i>	Select this option to enable the simple traffic policing style for the policy-class. The simple form of the police attribute uses a single data rate and burst size, resulting in two outcomes (conform and violate). After you select this option, configure the following policing criteria: • Color Mode The type of color policing used in DiffServ traffic conditioning. • Color Conform Class For color-aware policing, packets in this class are metered against both the committed information rate (CIR) and the peak information rate (PIR). The class definition used for policing color awareness is only allowed to contain a single, non-excluded class match condition identifying one of the supported comparison fields: <i>CoS</i>, <i>IP DSCP</i>, <i>IP Precedence</i>, or <i>Secondary CoS</i>. • Committed Rate (Kbps) The maximum allowed arrival rate of incoming packets for this class. • Committed Burst Size (Kbytes) The amount of conforming traffic allowed in a burst. • Conform Action The action taken on packets considered to be conforming (below the police rate). • Violate Action The action taken on packets considered to be non-conforming (above the police rate).
<i>Police Single Rate</i>	Select this option to enable the single-rate traffic policing style for the policy-class. The single-rate form of the police attribute uses a single data rate and two burst sizes, resulting in three outcomes (conform, exceed, and violate). After you select this option, configure the following policing criteria: • Color Mode The type of color policing used in DiffServ traffic conditioning. • Color Conform Class For color-aware policing, packets are metered against the committed information rate (CIR) and the peak information rate (PIR). The class definition used for policing color awareness is only allowed to contain a single, non-excluded class match condition identifying one of the supported comparison fields: <i>CoS</i>, <i>IP DSCP</i>, <i>IP Precedence</i>, or <i>Secondary CoS</i>. This field is available only if one or more classes that meets the color-awareness criteria exist. • Color Exceed Class For color-aware policing, packets are metered against the PIR only. • Committed Rate (Kbps) The maximum allowed arrival rate of incoming packets for this class. • Committed Burst Size (Kbytes) The amount of conforming traffic allowed in a burst. • Excess Burst Size (Kbytes) The amount of conforming traffic allowed to accumulate beyond the Committed Burst Size (Kbytes) value during longer-than-normal idle times. This value allows for occasional bursting. • Conform Action The action taken on packets considered to be conforming (below the police rate). • Exceed Action The action taken on packets that are considered to exceed the committed burst size but are within the excessive burst size. • Violate Action The action taken on packets considered to be non-conforming (above the police rate).
<i>Police Two Rate</i>	Select this option to enable the two-rate traffic policing style for the policy-class. The two-rate form of the police attribute uses two data rates and two burst sizes. Only the smaller of the two data rates is intended to be guaranteed. After you select this option, configure the following policing criteria: • Color Mode The type of color policing used in DiffServ traffic conditioning. • Color Conform Class For color-aware policing, packets are metered against the committed information rate (CIR) and the peak information rate (PIR). The class definition used for policing color awareness is only allowed to contain a single, non-excluded class match condition identifying one of the supported comparison fields: <i>CoS</i>, <i>IP DSCP</i>, <i>IP Precedence</i>, or <i>Secondary CoS</i>. This field is available only if one or more classes that meets the color-awareness criteria exist. • Color Exceed Class For color-aware policing, packets are metered against the PIR. • Committed Rate (Kbps) The maximum allowed arrival rate of incoming packets for this class. • Committed Burst Size (Kbytes) The amount of conforming traffic allowed in a burst. • Peak Rate (Kbps) The maximum information rate for the arrival of incoming packets for this class. • Excess Burst Size (Kbytes) The maximum size of the packet burst that can be accepted to maintain the Peak Rate (Kbps). • Conform Action The action taken on packets considered to be conforming (below the police rate). • Exceed Action The action taken on packets that are considered to exceed the committed burst size but are within the excessive burst size. • Violate Action The action taken on packets considered to be non-conforming (above police rate).
<i>Redirect Interface</i>	Select this option to force a classified traffic stream to the specified egress port (physical port or LAG). Use the Interface field to select the interface to which traffic is redirected.

After you select the policy to configure from the Policy menu, use the buttons to perform the following tasks:

- To add a class to the policy, click **Add Class**.
- To add attributes to a policy or to change the policy attributes, select the policy with the attributes to configure and click **Add Attribute**.
- To remove the most recently associated class from the selected policy, click **Remove Last Class**.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Diffserv Service Summary

Use the *Diffserv Service Summary* page to add DiffServ policies to interfaces, remove policies from interfaces, and edit policy-interface mappings.

To display the page, click **QoS > Diffserv > Service Summary** in the navigation menu.



DiffServ Service Summary

Diffserv Service Summary Fields

Field	Description
<i>Interface</i>	The interface associated with the rest of the data in the row. Only interfaces that have an associated policy are listed in the table.
<i>Direction</i>	The traffic flow direction to which the policy is applied: Inbound The policy is applied to traffic as it enters the interface. Outbound The policy is applied to traffic as it exits the interface.
<i>Status</i>	The status of the policy on the interface. A policy is <i>Up</i> if DiffServ is globally enabled, and if the interface is administratively enabled and has a link. Otherwise, the status is <i>Down</i> .
<i>Policy</i>	The DiffServ policy associated with the interface.
<i>Configure Service</i> dialog box fields – When you click Add or Edit , this dialog box opens and allows you to configure DiffServ interface policies. Specifying <i>None</i> for a policy has no effect when adding or editing interface policies. To remove an interface policy mapping, click Remove on the parent page. The following information describes the fields in this window.	
<i>Interface</i>	Select an interface to associate with a policy.
<i>Policy In</i>	The menu lists all policies configured with a type of <i>In</i> . Select the policy to apply to traffic as it enters the interface.
<i>Policy Out</i>	The menu lists all policies configured with a type of <i>Out</i> . Select the policy to apply to traffic as it exits the interface.

Use the buttons to perform the following tasks:

- To add a policy to an interface, click **Add**, configure the fields, and click **Submit** to apply the changes.
- To edit a configured interface-policy association, select the entry to modify and click **Edit**. Configure the fields, and click **Submit** to apply the changes.
- To remove one or more configured interface-policy associations, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.
- Click **Refresh** to refresh the page with the most current data from the switch.

To retain the changes across the switch's next power cycle, click **System > Configuration Storage > Save**.

Diffserv Service Performance Statistics

This page displays service-level statistical information for all interfaces in the system to which a DiffServ policy has been attached.

To display the *Diffserv Service Performance Statistics* page, click **QoS > Diffserv > Service Statistics** in the navigation menu.



Diffserv Service Performance Statistics

Diffserv Service Performance Statistics Fields

Field	Description
<i>Interface</i>	The interface associated with the rest of the data in the row. The table displays all interfaces that have a DiffServ policy currently attached in a traffic flow direction.
<i>Direction</i>	The traffic flow direction to which the policy is applied: In The policy is applied to traffic as it enters the interface. Out The policy is applied to traffic as it exits the interface.
<i>Status</i>	The operational status of this service interface, either <i>Up</i> or <i>Down</i> .
<i>Octets Offered</i>	The total number of octets offered to all class instances in this service policy before their defined DiffServ treatment is applied. This is the overall count per-interface, per-direction.
<i>Octets Discarded</i>	The total number of octets discarded for all class instances in this service policy for any reason due to DiffServ treatment. This is the overall count per-interface, per-direction.
<i>Octets Sent</i>	The total number of octets forwarded for all class instances in this service policy after their defined DiffServ treatments were applied. In this case, forwarding means the traffic stream was passed to the next functional element in the data path, such as the switching or routing function of an outbound link transmission element. This is the overall count per-interface, per-direction.

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

Diffserv Policy Performance Statistics

This page displays class-oriented statistical information for the policy, which is specified by the interface and direction.

To display the *Diffserv Policy Performance Statistics* page, click **QoS > Diffserv > Policy Statistics** in the navigation menu.



Diffserv Policy Performance Statistics

Diffserv Policy Performance Statistics Fields

Field	Description
<i>Interface</i>	The interface associated with the rest of the data in the row. The table displays all interfaces that have a DiffServ policy currently attached in a traffic flow direction.
<i>Direction</i>	The traffic flow direction to which the policy is applied: In The policy is applied to traffic as it enters the interface. Out The policy is applied to traffic as it exits the interface.
<i>Policy</i>	The name of the policy currently attached to the interface.
<i>Status</i>	The operational status of the policy currently attached to the interface.
<i>Class</i>	The DiffServ class currently defined for the attached policy.
<i>Packets Offered</i>	The total number of packets offered to all class instances in this service policy before their defined DiffServ treatment is applied. This is the overall count per-interface, per-direction.
<i>Packets Discarded</i>	The total number of packets discarded for all class instances in this service policy for any reason due to DiffServ treatment. This is the overall count per-interface, per-direction.

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

Appendix A: Configuration Examples

This appendix contains examples of how to configure selected features available in the EdgeSwitch software. Each example contains procedures on how to configure the feature by using the EdgeSwitch UI or CLI.

This appendix describes how to perform the following procedures:

- [“Configuring VLANs” on page 259](#)
- [“Configuring Multiple Spanning Tree Protocol” on page 262](#)
- [“Configuring VLAN Routing” on page 265](#)
- [“Configuring Policy-Based Routing” on page 267](#)
- [“Configuring 802.1X Network Access Control” on page 270](#)
- [“Configuring Differentiated Services for VoIP” on page 271](#)




Note: Each configuration example starts from a factory-default configuration unless otherwise noted.

Configuring VLANs

The diagram in this section shows a switch with four ports configured to handle the traffic for two VLANs. Port 0/2 handles traffic for both VLANs, while port 0/1 is a member of VLAN 2 only, and ports 0/3 and 0/4 are members of VLAN 3 only.

The following examples show how to create VLANs, assign ports to the VLANs, and assign a VLAN as the default VLAN to a port.




VLAN Example Network Diagram

Using the EdgeSwitch UI to Configure VLANs

1. Access the **Switching > VLAN > Status** page.
2. Click **Add** to create a new VLAN.

3. Type **2-3** in the *VLAN ID or Range* field, and then click **Submit**.



The 'Add VLAN' dialog box features a text input field labeled 'VLAN ID or Range'. Below the field, a note states: 'Enter VLAN ID in the range 2 to 4094. Use '-' to specify a range and ',' to separate VLAN IDs or VLAN ranges in the list.' At the bottom right, there are 'Submit' and 'Cancel' buttons.

Add VLAN Dialog Box

4. From the *Port Configuration* page, select **VLAN 2** from the *VLAN ID List*.
5. From the *Participation* column in the interface table, select **Include** for ports 0/1 and 0/2 to specify that these ports are members of VLAN 2.
6. Select the interface check box and click **Edit**. Select the **Tagging All** check box to specify that frames will always be transmitted tagged from ports that are members of VLAN 2.
7. Click **Submit**.
8. Select **VLAN 3** from the *VLAN ID and Name List*.
9. Select the **Participate** option in the *VLAN* field.
10. For ports 0/2, 0/3, and 0/4, select **Include** from the *Participation* menu to specify that these ports are members of VLAN 3.
11. Click **Submit**.
12. Go to the **Switching > VLAN > Port Summary** page.
13. In the *Interface* column, select **0/1** and click **Edit**.
14. In the *Acceptable Frame Type* field, select **Only Tagged** to specify that untagged frames will be rejected on receipt.
15. Click **Submit**.
16. In the *Interface* column, select **0/2** and click **Edit**.
17. In the *Port VLAN ID* field, enter **3** to assign VLAN 3 as the default VLAN for the port.
18. In the *Acceptable Frame Types* field, select **Admit All** to specify the untagged frames will be rejected on receipt.



The 'VLAN Port Configuration' page shows a table of ports with columns for 'Interface', 'VLAN', 'Participation', 'Acceptable Frame Type', and 'Acceptable Frame Types'. The table lists ports 0/1 through 0/24. Port 0/1 is selected, and its configuration is shown in a detail view at the bottom. The detail view includes fields for 'Port VLAN ID' (set to 3) and 'Acceptable Frame Types' (set to Admit All). Buttons for 'Submit' and 'Cancel' are at the bottom.

VLAN Port Configuration Page

19. Click **Submit**.

Using the CLI to Configure VLANs

1. Create VLAN 2 and VLAN 3.

```
(UBNT EdgeSwitch) #vlan database
vlan 2
vlan 3
exit
```

2. Assign ports 0/1 and 0/2 to VLAN2 and specify that untagged frames will be rejected on receipt.

```
(UBNT EdgeSwitch) #Config
interface 0/1
    vlan participation include 2
    vlan acceptframe vlanonly
    exit
interface 0/2
    vlan participation include 2
    vlan acceptframe all
```

3. While in interface config mode for port 0/2, assign VLAN3 as the default VLAN.

```
(UBNT EdgeSwitch) (Interface 0/2)#vlan pvid 3
exit
```

4. Specify that frames will always be transmitted tagged from ports that are members of VLAN 2.

```
(UBNT EdgeSwitch)(Config)#vlan port tagging all 2
exit
```

5. Assign the ports that will belong to VLAN 3.



Note: Port 0/2 belongs to both VLANs, and port 0/1 can never belong to VLAN 3.

```
(UBNT EdgeSwitch) #Config
interface 0/2
    vlan participation include 3
    exit
interface 0/3
    vlan participation include 3
    exit
interface 0/4
    vlan participation include 3
    exit
exit
```

6. Specify that untagged frames will be accepted on port 0/4.

```
(UBNT EdgeSwitch) #Config
interface 0/4
    vlan acceptframe all
    exit
exit
```

Configuring Multiple Spanning Tree Protocol

This example shows how to enable IEEE 802.1s Multiple Spanning Tree (MST) protocol on the switch and all of the ports and to set the bridge priority.

To make multiple switches be part of the same MSTP region, make sure the Force Protocol Version setting for all switches is IEEE 802.1s. Also, make sure the configuration name, digest key, and revision level are the same for all switches in the region.



Note: The digest key is generated based on the association of VLANs to different instances. To ensure the digest key is same, the mapping of VLAN to instance must be the same on each switch in the region. For example, if VLAN 10 is associated with instance 10 on one switch, you must associate VLAN 10 and instance 10 on the other switches.


Using the Web UI to Configure MSTP

1. Create VLANs 10 and 20.
 - a. Access the **Switching > VLAN > Status** page.
 - b. Click **Add** to create a VLAN.
 - c. Select the **VLAN ID-Individual** option and enter **10**.
 - d. Click **Submit**.
 - e. Repeat the steps to add VLAN 20.
2. Enable MSTP (IEEE 802.1s) on the switch and change the configuration name.
 - a. Changing the configuration name allows all the bridges that want to be part of the same region to join.
 - b. Go to the **Switching > Spanning Tree > Switch** page.
 - c. From the *Spanning Tree Admin Mode* field, select **Enable**.
 - d. In the *Configuration Name* field, enter **ubnt**.
 - e. Click **Submit**.

Spanning Tree Switch Configuration Page

3. Create two MST instances.
 - a. Go to the **Switching > Spanning Tree > MST** page.
 - b. From the *MST* page, click **Add**.
 - c. In the *MST ID* field, enter **10**.
 - d. Associate MST ID 10 with VLAN 10 and assign a bridge priority of **16384**.
 - e. Click **Submit**.

- f. Repeat the steps to create an MST instance with an ID of **20**.

The image shows a dialog box titled "Add MST Entry". It has three main sections. The first section is for "MST ID" with a text input field and a range "(1 to 4094)". The second section is for "Priority" with a text input field containing "8000" and a range "(0 to F000 hex)". The third section is for "# of Associated VLANs" and contains two sub-sections: "Available VLANs" and "Members". The "Available VLANs" list shows numbers 1 through 8. The "Members" list is empty. At the bottom right of the dialog box are "Submit" and "Cancel" buttons.

Add MST Entry Dialog Box

4. Use similar procedures to associate MST instance 20 to VLAN 20 and assign it a bridge priority value of 61440.

By using a lower priority for MST 20, MST 10 becomes the root bridge.

5. Force port 0/2 to be the root port for MST 20, which is the non-root bridge.
 - a. Go to the **Switching > Spanning Tree > MST** page.
 - b. From the *MST ID* menu, select **20**.
 - c. From the *Interface* menu, select **0/2**.
 - d. In the *Port Priority* field, enter **64**.
 - e. Click **Submit**.

Using the CLI to Configure MSTP

1. Create VLAN 10 and VLAN 20.

```
(UBNT EdgeSwitch) #vlan database
vlan 10
vlan 20
exit
```

2. Enable spanning tree globally.

```
(UBNT EdgeSwitch) #config
spanning-tree
```

3. Create MST instances 10 and 20.

```
spanning-tree mst instance 10
spanning-tree mst instance 20
```

4. Associate MST instance 10 to VLAN 10 and MST instance 20 to VLAN 20.

```
spanning-tree mst vlan 10 10
spanning tree mst vlan 20 20
```

5. Change the name so that all the bridges that want to be part of the same region can form the region.

```
spanning-tree configuration name ubnt
```

6. Make the MST ID 10 bridge the root bridge by lowering the priority.

```
spanning-tree mst priority 10 16384
```

7. Change the priority of MST ID 20 to ensure the other bridge is the root bridge.

```
spanning-tree mst priority 20 61440
```

8. Enable STP on interface 0/1.

```
interface 0/1
spanning-tree port mode
exit
```

9. Enable STP on interface 0/2.

```
interface 0/2
spanning-tree port mode
```


10. On the non-root bridge, change the priority to force port 0/2 to be the root port.

```
spanning-tree mst 20 port-priority 64
exit
```

Configuring VLAN Routing

This section provides an example of how to configure the EdgeSwitch software to support VLAN routing. The configuration of the VLAN router port is similar to that of a physical port. The main difference is that, after the VLAN has been created, you must use the `show ip vlan` command to determine the VLAN's interface ID so that you can use it in the router configuration commands.

The diagram in this section shows a Layer-3 switch configured for port routing. It connects two VLANs, with two ports participating in one VLAN, and one port in the other. The script shows the commands you would use to configure the EdgeSwitch software to provide the VLAN routing support shown in the diagram.



VLAN Routing Example Network Diagram

Using the CLI to Configure VLAN Routing

1. Create VLAN 10 and VLAN 20.

```
(UBNT EdgeSwitch) #vlan database
vlan 10
vlan 20
exit
```

2. Configure ports 0/1, 0/2 as members of VLAN 10 and specify that untagged frames received on these ports will be assigned to VLAN 10.

```
config
interface 0/1
vlan participation include 10
vlan pvid 10
exit
interface 0/2
vlan participation include 10
vlan pvid 10
exit
```

3. Configure port 0/3 as a member of VLAN 20 and specify that untagged frames received on these ports will be assigned to VLAN 20.

```
interface 0/3
  vlan participation include 20
  vlan pvid 20
  exit
exit
```

4. Specify that all frames transmitted for VLANs 10 and 20 will be tagged.

```
config
  vlan port tagging all 10
  vlan port tagging all 20
  exit
```

5. Enable routing for the VLANs:

```
(UBNT EdgeSwitch) #vlan database
  vlan routing 10
  vlan routing 20
  exit
```

6. View the logical interface IDs assigned to the VLAN routing interfaces.

```
(UBNT EdgeSwitch) #show ip vlan
```

MAC Address used by Routing VLANs: 00:00:AA:12:65:12

VLAN ID	Logical Interface	IP Address	Subnet Mask
10	4/1	0.0.0.0	0.0.0.0
20	4/2	0.0.0.0	0.0.0.0

As the output shows, VLAN 10 is assigned ID 4/1 and VLAN 20 is assigned ID 4/2.

7. Enable routing for the switch:

```
config
  ip routing
  exit
```

8. Configure the IP addresses and subnet masks for the virtual router ports.

```
config
  interface 4/1
    ip address 192.150.3.1 255.255.255.0
  exit
  interface 4/2
    ip address 192.150.4.1 255.255.255.0
  exit
exit
```

Configuring Policy-Based Routing

In present-day networks, network administrators who manage organizations should be provided with a choice for implementing packet forwarding/routing according to the organization's policies. Policy-Based Routing (PBR) is a feature that fits this purpose. PBR provides a flexible mechanism to implement solutions in cases where organizational constraints dictate that traffic be routed through specific network paths.


Configuring PBR involves configuring a route-map with match and set commands and then applying the corresponding route-map to the interface.

Policy-Based Routing is applied to inbound traffic on physical routing/VLAN routing interfaces. Enabling the feature causes the router to analyze all packets incoming on the interface using a route-map configured for that purpose. One interface can only have one route-map tag, but an administrator can have multiple route-map entries with different sequence numbers. These entries are evaluated in sequence number order until the first match. If there is no match, the packets are routed as usual.

Configuring Policy-Based Routing Using the CLI

In the following configuration example, we have a Layer3 Switch/Router with four VLAN routing interfaces – VLAN 10, VLAN 20, VLAN 30, and VLAN 40. Each of these interfaces is connected to an L2 network.

- Physical Interface 0/2 – Member of VLAN 10
- Physical Interface 0/4 – Member of VLAN 20
- Physical Interface 0/22 – Member of VLAN 30
- Physical Interface 0/24 – Member of VLAN 40



Policy-Based Routing Example

In this example, the procedure to configure policy route traffic from VLAN routing interface 10 to VLAN routing interface 30 is shown in the diagram above. Traffic sent to VLAN Interface 10 is destined for VLAN Interface 20. In order to override the traditional destination routing and send the same traffic to VLAN Interface 30, use the following procedure.

1. Create VLANs 10, 20, 30, 40, and enable routing on these VLANs.

```
(UBNT EdgeSwitch) #vlan database
vlan 10,20,30,40
vlan routing 10 1
vlan routing 20 2
vlan routing 30 3
vlan routing 40 4
exit
```

2. Add physical ports to the VLANs and configure PVID on the corresponding interfaces.

```
config
interface 0/2
vlan pvid 10
vlan participation exclude 1
vlan participation include 10
exit
interface 0/4
vlan pvid 20
vlan participation exclude 1
vlan participation include 20
exit
interface 0/22
vlan pvid 30
vlan participation exclude 1
vlan participation include 30
exit
interface 0/24
vlan pvid 40
vlan participation exclude 1
vlan participation include 40
exit
exit
```

3. Enable routing on each VLAN interface and assign an IP address.

```
config
interface vlan 10
routing
ip address 1.1.1.1 255.255.255.0
exit
interface vlan 20
routing
ip address 2.2.2.1 255.255.255.0
exit
interface vlan 30
routing
ip address 3.3.3.1 255.255.255.0
exit
interface vlan 40
routing
ip address 4.4.4.3 255.255.255.0
exit
```

4. Enable IP Routing (Global configuration).

```
config
    ip routing
exit
```

After this step, if traffic with the following characteristics is sent, it will be routed from VLAN routing interface 10 to VLAN routing interface 20.

Source IP: 1.1.1.2
Destination IP: 2.2.2.2

In order to policy route such traffic to VLAN routing interface 30, continue with the following steps:

5. Create an access-list matching incoming traffic.

```
config
    access-list 1 permit 1.1.1.2 0.0.0.255
exit
```

6. Create a route-map and add match/set terms to the route-map.

```
configure
    route-map pbr_test permit 10
    match ip address 1
    set ip next-hop 3.3.3.3
    exit
exit
```

7. Assign a route-map to VLAN routing interface 10.

```
config
    interface vlan 10
    ip policy pbr_test
    exit
exit
```

After this step, traffic mentioned in the diagram **“Policy-Based Routing Example” on page 267** is policy-routed to VLAN interface 30. Counters are incremented in the “show route-map” command indicating that traffic is being policy routed.

8. Run the show command.

```
(UBNT EdgeSwitch) #show route-map pbr_test
route-map pbr_test permit 10
```

Match clauses:

```
ip address (access-lists) : 1
```


Set clauses:

```
ip next-hop 3.3.3.3
```

Policy routing matches: 19922869 packets, 1275063872 bytes

Configuring 802.1X Network Access Control

This example configures a single RADIUS server used for authentication and accounting at 10.10.10.10. The shared secret is configured to be secret. The switch is configured to require that the 802.1X access method is through a RADIUS server. IEEE 802.1X port-based access control is enabled for the system, and interface 0/1 is configured to be in force-authorized mode because this is where the RADIUS server and protected network resources are located.



If a user, or supplicant, attempts to communicate via the switch on any interface except interface 0/1, the system challenges the supplicant for login credentials. The system encrypts the provided information and transmits it to the RADIUS server. If the RADIUS server grants access, the system sets the 802.1X port state of the interface to authorized, and the supplicant is able to access network resources.

Using the CLI to Configure 802.1X Port-Based Access Control

1. Configure the RADIUS authentication server IP address.

```
(UBNT EdgeSwitch) #config
radius server host auth 10.10.10.10
```

2. Configure the RADIUS authentication server secret.

```
radius server key auth 10.10.10.10
secret
secret
```

3. Configure the RADIUS accounting server IP address.

```
radius server host acct 10.10.10.10
```

4. Configure the RADIUS accounting server secret.

```
radius server key acct 10.10.10.10
secret
secret
```

5. Enable RADIUS accounting mode.

```
radius accounting mode
```

6. Set IEEE 802.1X to use RADIUS as the AAA method.

```
aaa authentication dot1x default radius
```

7. Enable 802.1X authentication on the switch.


```
dot1x system-auth-control
```

- Set the 802.1X mode for port 0/1 to Force Authorized.

```
interface 0/1
  dot1x port-control force-authorized
exit
```

Configuring Differentiated Services for VoIP

One of the most valuable uses of DiffServ is to support Voice over IP (VoIP). VoIP traffic is inherently time-sensitive: for a network to provide acceptable service, a guaranteed transmission rate is vital. This example shows one way to provide the necessary quality of service: how to set up a class for UDP traffic, have that traffic marked on the inbound side, and then expedite the traffic on the outbound side. The configuration script is for Router 1 in the accompanying diagram: a similar script should be applied to Router 2.



Using the CLI to Configure DiffServ VoIP Support

- Enter Global Config mode. Set queue 5 on all ports to use strict priority mode. This queue shall be used for all VoIP packets. Activate DiffServ for the switch.

```
(UBNT EdgeSwitch) #config
cos-queue strict 5
diffserv
```

- Create a DiffServ classifier named 'class_voip' and define a single match criterion to detect UDP packets. The class type match-all indicates that all match criteria defined for the class must be satisfied in order for a packet to be considered a match.

```
class-map match-all class_voip
  match protocol udp
exit
```

3. Create a second DiffServ classifier named 'class_ef' and define a single match criterion to detect a DiffServ code point (DSCP) of 'EF' (expedited forwarding). This handles incoming traffic that was previously marked as expedited elsewhere in the network.

```
class-map match-all class_ef
  match ip dscp ef
  exit
```

4. Create a DiffServ policy for inbound traffic named 'pol_voip', and then add the previously created classes 'class_ef' and 'class_voip' as instances within this policy.

This policy handles incoming packets already marked with a DSCP value of 'EF' (per 'class_ef' definition), or marks UDP packets per the 'class_voip' definition) with a DSCP value of 'EF'. In each case, the matching packets are assigned internally to use queue 5 of the egress port to which they are forwarded.

```
policy-map pol_voip in
  class class_ef
    assign-queue 5
  exit
  class class_voip
    mark ip-dscp ef
    assign-queue 5
  exit
exit
```

Attach the defined policy to an inbound service interface.

```
interface 0/2
  service-policy in pol_voip
  exit
exit
```

Appendix B: Contact Information

Ubiquiti Networks Support

Ubiquiti Support Engineers are located around the world and are dedicated to helping customers resolve software, hardware compatibility, or field issues as quickly as possible. We strive to respond to support inquiries within a 24-hour period.

Online Resources

Support: support.ubnt.com

Community: community.ubnt.com

Downloads: downloads.ubnt.com



Ubiquiti Networks, Inc.
2580 Orchard Parkway
San Jose, CA 95131
www.ubnt.com

©2014 Ubiquiti Networks, Inc. All rights reserved. Ubiquiti, Ubiquiti Networks, the Ubiquiti U logo, the Ubiquiti beam logo, EdgeMAX, and EdgeSwitch are trademarks or registered trademarks of Ubiquiti Networks, Inc. in the United States and in other countries. All other trademarks are the property of their respective owners.