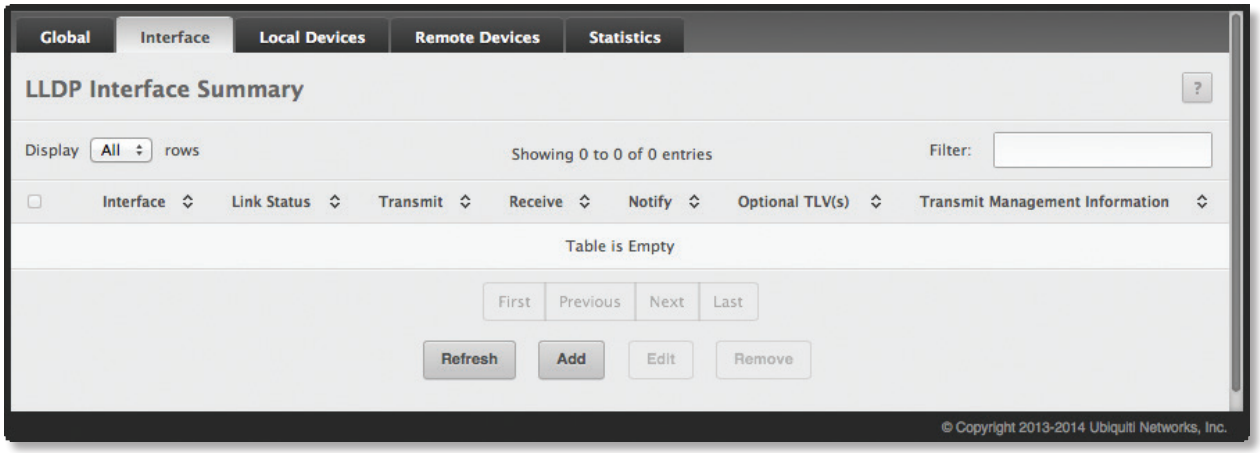
LLDP Interface Configuration

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
Interface	The interface associated with the rest of the data in the row. Only interfaces that have at least one LLDP setting enabled appear in the table. In the *Add LLDP Interface* window, use this field to select the interface with the LLDP settings to configure. In the *Edit LLDP Interface* window, this field identifies the interface that is being configured.
Link Status	The link status of the interface: *Up* or *Down*. An interface that is down does not forward traffic.
Transmit	The LLDP advertise (transmit) mode on the interface. If the transmit mode is enabled, the interface sends LLDP Data Units (LLDPDUs) that advertise the mandatory TLVs and any optional TLVs that are enabled.
Receive	The LLDP receive mode on the interface. If the receive mode is enabled, the device can receive LLDPDUs from other devices.
Notify	The LLDP remote data change notification status on the interface. If the notify mode is enabled, the interface sends SNMP notifications when a link partner device is added or removed.
Optional TLV(s)	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**.
Transmit Management Information	Select the check box to enable the transmission of management address instance. Clear the check box to disable management information transmission. The default is *Disabled*.
Add/Edit LLDP Interface dialog box – Click **Add** or **Edit** to open a dialog box and configure the LLDP settings for an interface:	
Port Description	Select this option to include the user-configured port description in the LLDPDU the interface transmits.
System Name	Select this option to include the user-configured system name in the LLDPDU the interface transmits. The system name, configured on the *System Description* page, is the SNMP server name for the device.
System Description	Select this option to include a description of the device in the LLDPDU the interface transmits. The description includes information about the product model and platform.
System Capabilities	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:

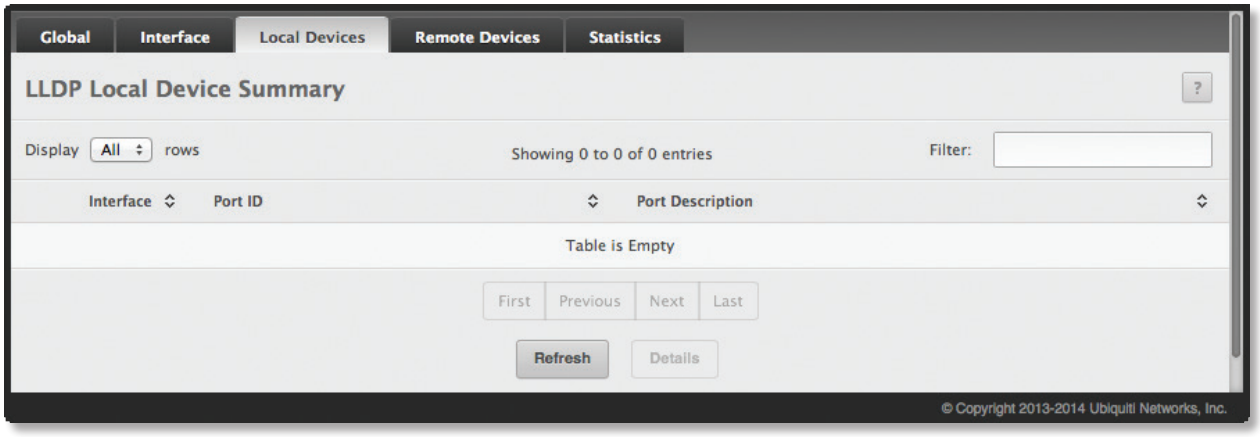
- To configure LLDP settings on an interface that does not have any LLDP settings enabled, click **Add**, configure the settings as needed, and click **Submit** to apply the new settings to the system.
- To change the LLDP settings for an interface in the table, select the entry to update, click **Edit**, 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

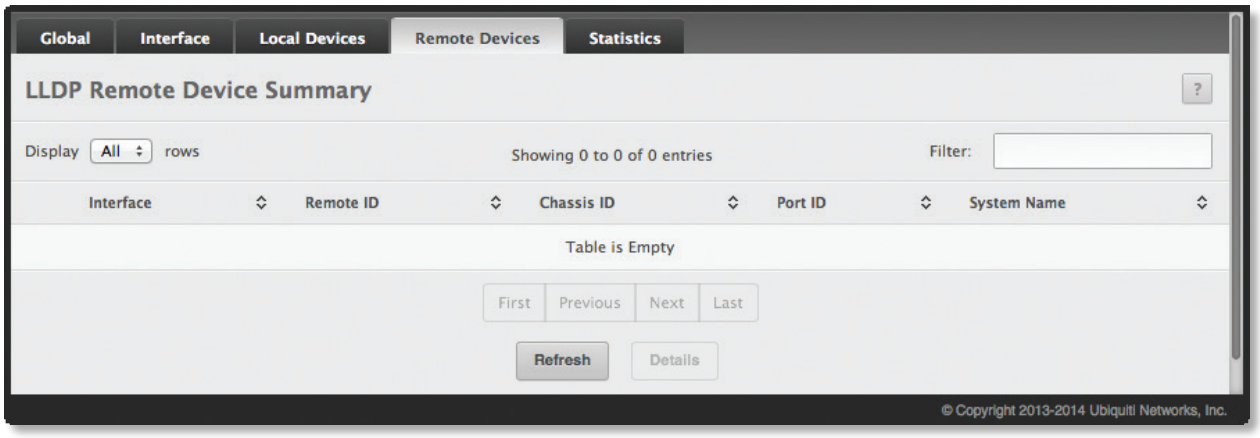
Field	Description
Interface	The interface associated with the rest of the LLDP - 802.1AB data in the row. When viewing the details for an interface, this field identifies the interface that is being viewed.
Port ID	The port identifier, which is the physical address associated with the interface.
Port Description	A description of the port. An administrator can configure this information on the *Port Description* page.
Click **Details** to display the following additional information about the data the interface transmits in its LLDPDUs:	
Chassis ID Subtype	The type of information used to identify the device in the *Chassis ID* field.
Chassis ID	The hardware platform identifier for the device.
Port ID Subtype	The type of information used to identify the interface in the *Port ID* field.
System Name	The user-configured system name for the device. The system name is configured on the *System Description* page and is the SNMP server name for the device.
System Description	The device description, which includes information about the product model and platform.
System Capabilities Supported	The primary function(s) the device supports.
System Capabilities Enabled	The primary function(s) the device supports that are enabled.
Management Address	The physical address associated with the management interface of the device.
Management Address Type	The protocol type or standard associated with the management address.

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

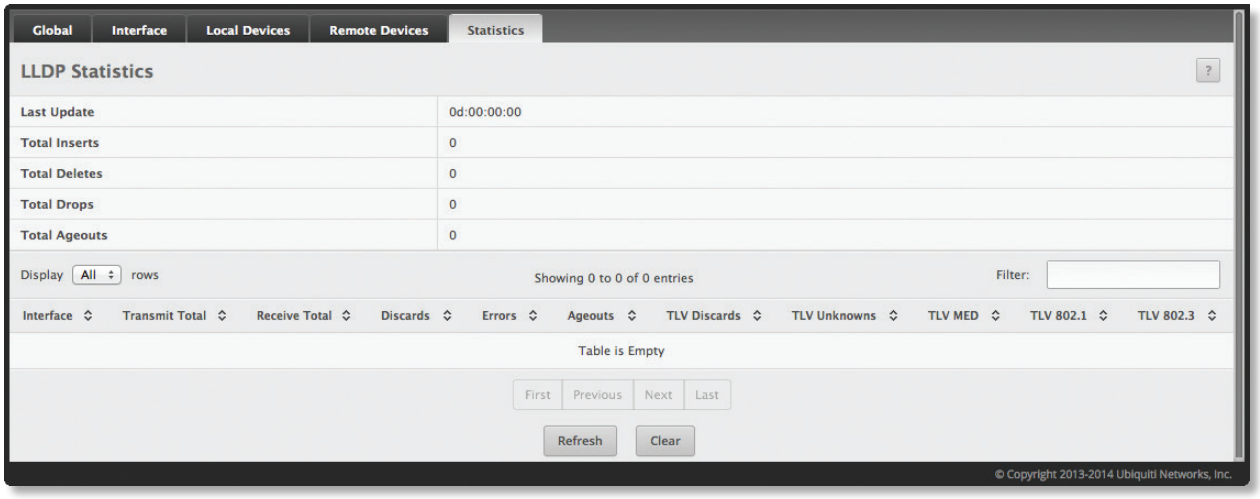
Field	Description
Interface	The local interface that is enabled to receive LLDPDUs from remote devices.
Remote ID	The client identifier assigned to the remote system that sent the LLDPDU.
Chassis ID	The information the remote device sent as the Chassis ID TVL. This identifies the hardware platform for the remote system.
Port ID	The port on the remote system that transmitted the LLDP data.
System Name	The system name configured on the remote device.
Click **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.	
Chassis ID Subtype	The type of information used to identify the device in the *Chassis ID* field.
Port ID Subtype	The type of information used to identify the interface in the *Port ID* field.
System Description	The device description, which includes information about the product model and platform.
Port Description	The description of the port on the remote device that transmitted the LLDP data.
System Capabilities Supported	The primary function(s) the remote system supports. The possible capabilities include *Other*, *Repeater*, *Bridge*, *WLAN AP*, *Router*, *Telephone*, *DOCSIS cable device*, and *Station*.
System Capabilities Enabled	The primary function(s) of the remote system that are both supported and enabled. The possible capabilities include *Other*, *Repeater*, *Bridge*, *WLAN AP*, *Router*, *Telephone*, *DOCSIS cable device*, and *Station*.
Time To Live	The number of seconds the local device should consider the LLDP data it received from the remote system to be valid.

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
Last Update	Displays the time when an entry was created, modified, or deleted in the tables associated with the remote systems.
Total Inserts	Displays the number of times a complete set of information advertised by a particular MAC Service Access Point (MSAP) has been inserted into the tables associated with the remote systems.
Total Deletes	Displays the number of times a complete set of information advertised by a particular MAC Service Access Point (MSAP) has been deleted from the tables associated with the remote systems.
Total Drops	Displays the number of times a complete set of information advertised by a particular MAC Service Access Point (MSAP) could not be entered into tables associated with the remote systems because of insufficient resources.
Total Ageouts	Displays the number of times a complete set of information advertised by a particular MAC Service Access Point (MSAP) has been deleted from tables associated with the remote systems because the information timelines interval has expired.
Interface	Identifies the interfaces.
Transmit Total	Displays the total number of LLDP frames transmitted by the LLDP agent on the corresponding port.
Receive Total	Displays the total number of valid LLDP frames received by the LLDP agent on the corresponding port, while the LLDP agent is enabled.
Discards	Displays the number of LLDP TLVs discarded for any reason by the LLDP agent on the corresponding port.
Errors	Displays the number of invalid LLDP frames received by the LLDP agent on the corresponding port, while the LLDP agent is enabled.
Ageouts	Displays the number of age-outs that occurred on a given port. An age-out is the number of times the complete set of information advertised by a particular MAC Service Access Point (MSAP) has been deleted from tables associated with remote entries because the information timeliness interval had expired.
TLV Discards	Displays the number of LLDP TLVs (Type, Length, Value sets) discarded for any reason by the LLDP agent on the corresponding port.
TLV Unknowns	Displays the number of LLDP TLVs received on the local ports which were not recognized by the LLDP agent on the corresponding port.
TLV MED	Displays the total number of LLDP-MED TLVs received on the local ports.
TLV 802.1	Displays the total number of LLDP TLVs received on the local ports which are of type 802.1.
TLV 802.3	Displays the total number of LLDP TLVs received on the local ports which are of type 802.3.

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
Fast Start Repeat Count	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*.
Device Class	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:

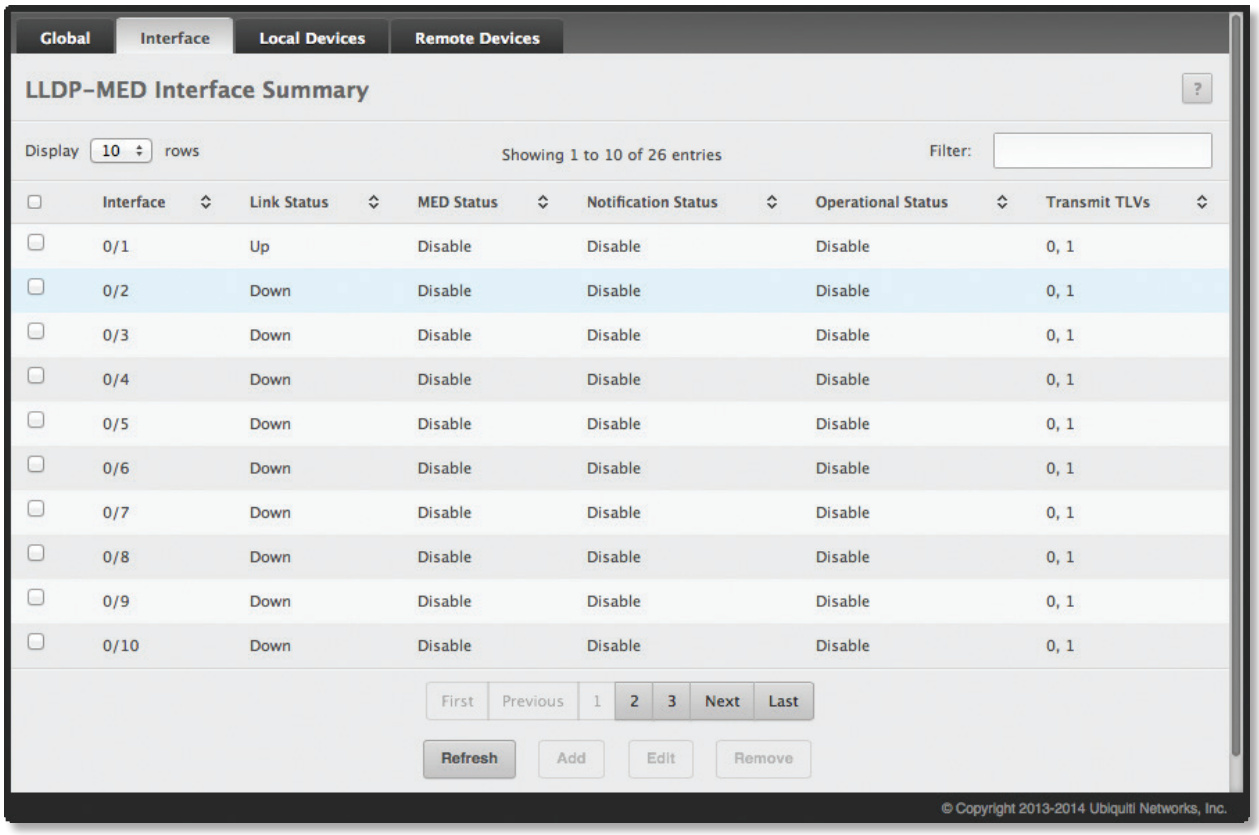
- If you make any changes on this page, click **Submit** to apply the changes.
- Click **Refresh** to refresh the page with the most current data from the switch.

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.



LLDP-MED Interface Summary

LLDP-MED Interface Summary Fields

Field	Description
Interface	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.
Link Status	The link status of the interface, which is either *Up* or *Down*. An interface that is down does not forward traffic.
MED Status	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.
Notification Status	Indicates whether LLDP-MED topology change notifications are enabled or disabled on the interface.
Operational Status	Indicates whether the interface will transmit TLVs.
Transmit TLVs	The LLDP-MED TLV(s) that the interface transmits: • **0** Capabilities • **1** Network Policy

Use the buttons to perform the following tasks:

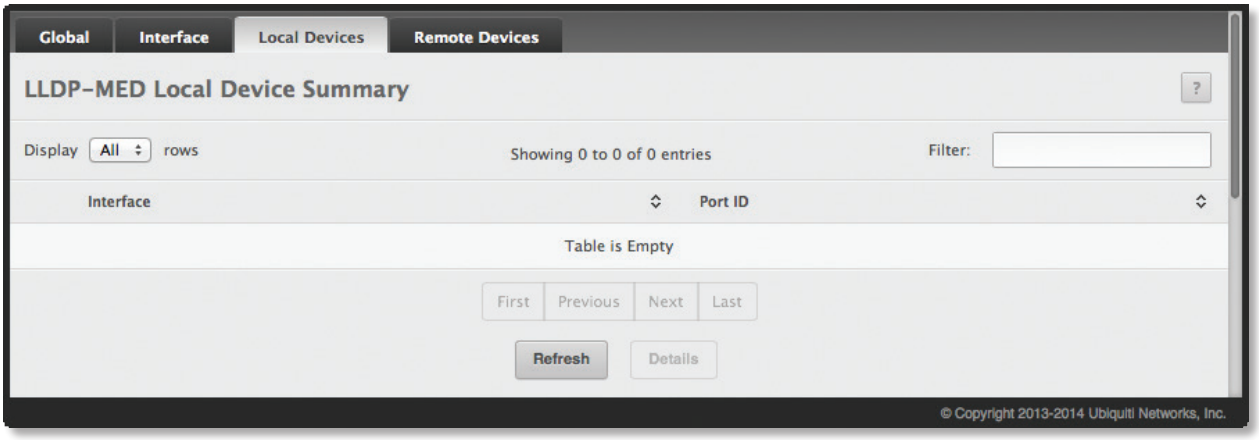
- To add an entry to the table, click **Add**, configure the fields, and click **Submit** to apply the changes.
- To change an entry, select the entry from table, and click **Edit**. When you have completed the changes, click **Submit** to apply the changes.
- To remove an entry, select it from the table, 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 LLPD-MED information advertised on the selected local interface.

To display this page, click **Switching** > **LLDP-MED** > **Local Devices** in the navigation menu.



LLPD-MED Local Device Summary

LLPD-MED Local Device Summary Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When viewing LLDP-MED details for an interface, this field identifies the interface that is being viewed.
Port ID	The MAC address of the interface. This is the MAC address that is advertised in LLDP-MED PDUs.
Network Policy Information – When you click **Details**, the *LLDP-MED Local Device Information* dialog box opens and shows the following detailed information about the LLDP-MED information the selected interface transmits.	
Media Application Type	The media application type transmitted in the TLV. The application types are *unknown*, *voicesignalling*, *guestvoice*, *guestvoicesignalling*, *softphonevoice*, *videoconferencing*, *streamingvideo*, and *videosignalling*. Each application type that is transmitted has the VLAN ID, priority, DSCP, tagged bit status, and unknown bit status. A port may transmit one or many such application types. This information is displayed only when a network policy TLV has been transmitted.
VLAN ID	The VLAN ID associated with a particular policy type.
Priority	The user priority associated with a particular policy type.
DSCP	The DSCP value associated with a particular policy type.
Unknown Bit Status	The unknown bit associated with a particular policy type.
Tagged Bit Status	Identifies whether the network policy is defined for tagged or untagged VLANs.
Location Information:	
Sub Type	The type of location information: • **Coordinate Based** The location map coordinates (latitude, longitude and altitude) of the device. • **Civic Address** The civic or street address location of the device. • **ELIN** The Emergency Call Service (ECS) Emergency Location Identification Number (ELIN) of the device.
Information	This column displays the information related to the coordinates, civic address, and ELIN for the device.

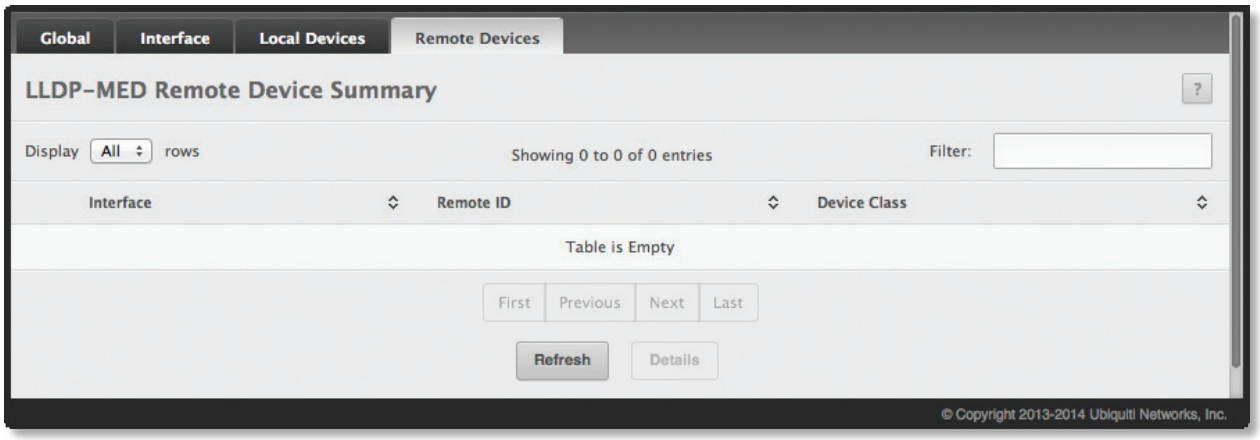
Use the buttons to perform the following tasks:

- 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
Interface	The local interface that has received LLDP-MED data units from remote devices.
Remote ID	The client identifier assigned to the remote system that sent the LLDP-MED data unit.
Capability Information:	
Supported Capabilities	The supported capabilities that were received in the MED TLV on this interface.
Enabled Capabilities	The supported capabilities on the remote device that are also enabled.
Device Class	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.
Network Policy Information:	
Media Application Type	The media application type received in the TLV from the remote device. The application types are *unknown, voicesignalling, guestvoice, guestvoicesignalling, softphonevoice, videoconferencing, streamingvideo*, and *videosignalling*. Each application type that is transmitted has the VLAN ID, priority, DSCP, tagged bit status, and unknown bit status. The port on the remote device may transmit one or many such application types. This information is displayed only when a network policy TLV has been received.
VLAN ID	The VLAN ID associated with a particular policy type.
Priority	The user priority associated with a particular policy type.
DSCP	The DSCP value associated with a particular policy type.
Unknown Bit Status	The unknown bit associated with a particular policy type.
Tagged Bit Status	Identifies whether the network policy is defined for tagged or untagged VLANs.
Inventory Information:	
Hardware Revision	The hardware version advertised by the remote device.
Firmware Revision	The firmware version advertised by the remote device.
Software Revision	The software version advertised by the remote device.
Serial Number	The serial number advertised by the remote device.

LLDP Remote Device Information Fields (Continued)

Field	Description
Manufacturer Name	The name of the system manufacturer advertised by the remote device.
Model Name	The name of the system model advertised by the remote device.
Asset ID	The system asset ID advertised by the remote device.
Location Information:	
Sub Type	The type of location information advertised by the remote device.
Information	The text description of the location information included in the subtype.
Extended PoE	Indicates whether the remote device is advertised as a PoE device.
Device Type	If the remote device is a PoE device, this field identifies the PoE device type of the remote device connected to this port.

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.

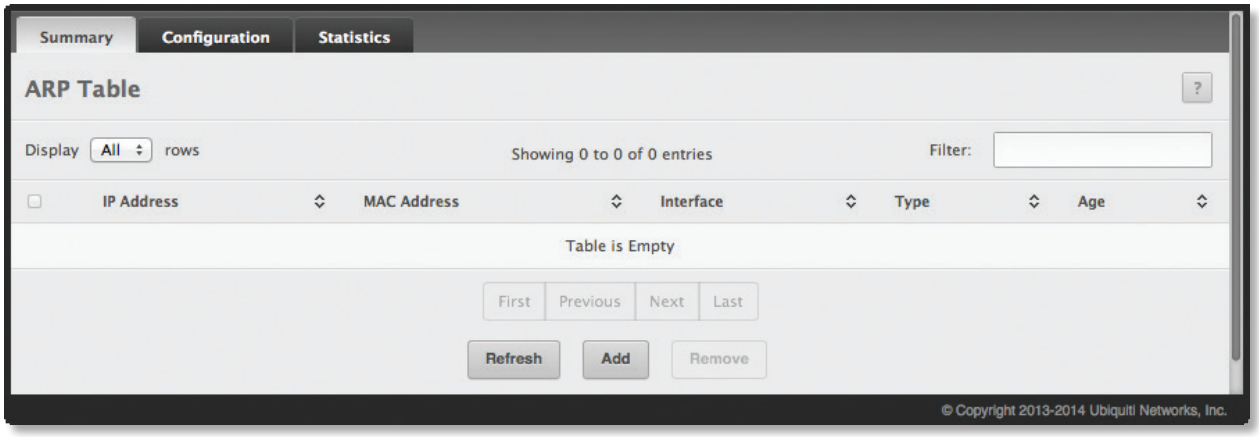
The **Routing** > **ARP Table** submenu contains links to the following UI pages that configure and display ARP-related details:

- **"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
IP Address	The IP address of a network host on a subnet attached to one of the device's routing interfaces. When adding a static ARP entry, specify the IP address for the entry after you click **Add**.
MAC Address	The unicast MAC address (hardware address) associated with the network host. When adding a static ARP entry, specify the MAC address to associate with the IP address in the entry.
Interface	The routing interface associated with the ARP entry. The network host is associated with the device through this interface.
Type	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
Age	The age of the entry since it was last learned or refreshed. This value is specified for *Dynamic* or *Gateway* entry types only (it is left blank for all other entry types).

Use the buttons to perform the following tasks:

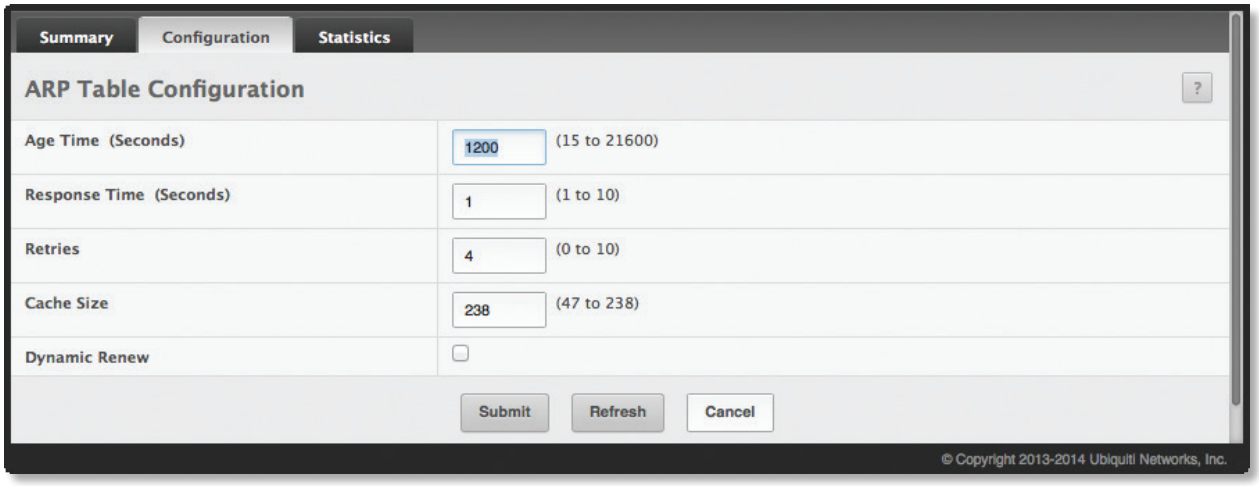
- To add a static ARP entry, click **Add**. In the *Add Static ARP Entry* dialog box, configure the information for the new ARP entry, and click **Submit** to apply the changes and save the entry to the ARP table.
- To delete one or more ARP entries, select each entry to delete and click **Remove**. 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
Age Time	The amount of time, in seconds, that a dynamic ARP entry remains in the ARP table before aging out.
Response Time	The amount of time, in seconds, that the device waits for an ARP response to an ARP request that it sends.
Retries	The maximum number of times an ARP request will be retried after an ARP response is not received. The number includes the initial ARP request.
Cache Size	The maximum number of entries allowed in the ARP table. This number includes all static and dynamic ARP entries.
Dynamic Renew	When selected, this option allows the ARP component to automatically attempt to renew dynamic ARP entries when they age out.

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

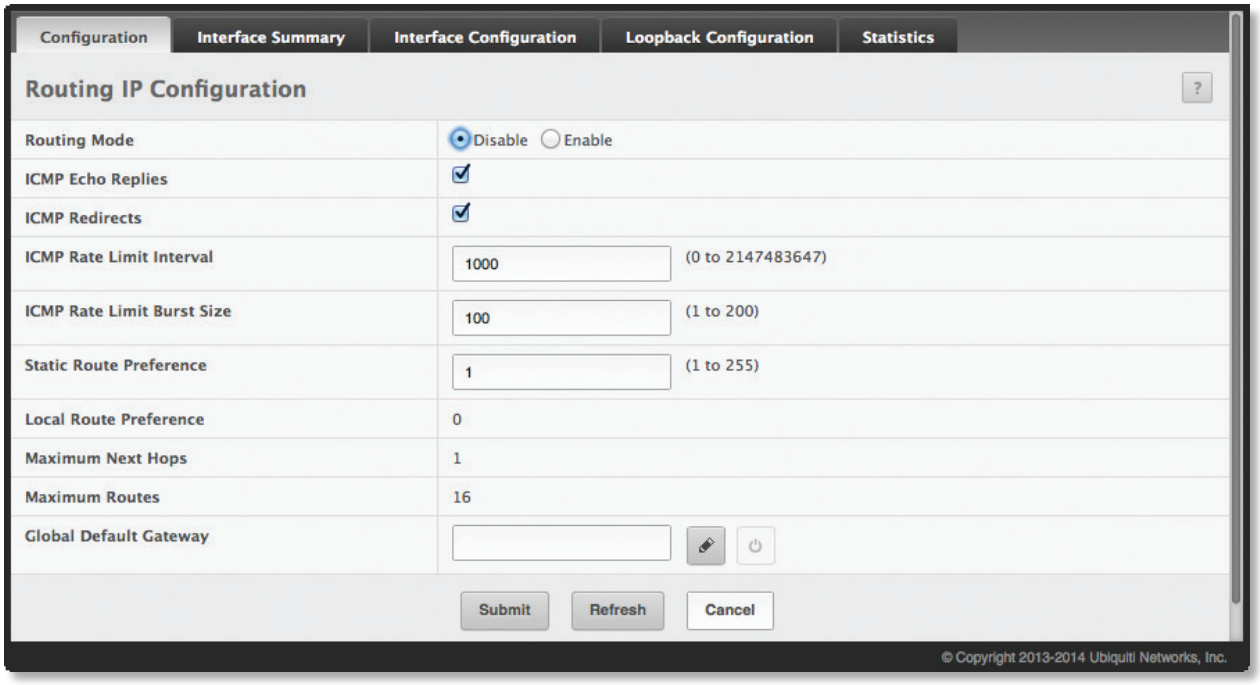
The **Routing** > **IP** folder contains links to the following UI pages that configure and display IP routing data:

- **"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
Routing Mode	The administrative mode of routing on the device. The options are as follows: • **Enable** The device can act as a Layer-3 device by routing packets between interfaces configured for IP routing. • **Disable** The device acts as a Layer-2 bridge and switches traffic between interfaces. The device does not perform any internetwork routing.
ICMP Echo Replies	Select *Enable* or *Disable* from the drop-down menu. If you select *Enable*, then only the router can send ECHO replies. By default, ICMP Echo Replies are sent for echo requests.
ICMP Redirects	Select this option to allow the device to send ICMP Redirect messages to hosts. An ICMP Redirect message notifies a host when a better route to a particular destination is available on the network segment.
ICMP Rate Limit Interval	To control the ICMP error packets, you can specify the number of ICMP error packets that are allowed per burst interval. By default, the rate limit is *100* packets per second, 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
ICMP Rate Limit Burst Size	To control the ICMP error packets, you can specify the number of ICMP error packets that are allowed per burst interval. By default, the burst size is *100* packets. When the burst interval is zero, then configuring this field is not a valid option. The valid burst size range is *1* to *200*.
Static Route Preference	The default distance (preference) for static routes. Lower route-distance values are preferred when determining the best route. The value configured for *Static Route Preference* is used when using the CLI to configure a static route and no preference is specified. Changing the *Static Route Preference* does not update the preference of existing static routes.
Local Route Preference	The default distance (preference) for local routes.
Maximum Next Hops	The maximum number of hops supported by the switch. This is a read-only value.
Maximum Routes	The maximum number of routes (routing table size) supported by the switch.
Global Default Gateway	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:

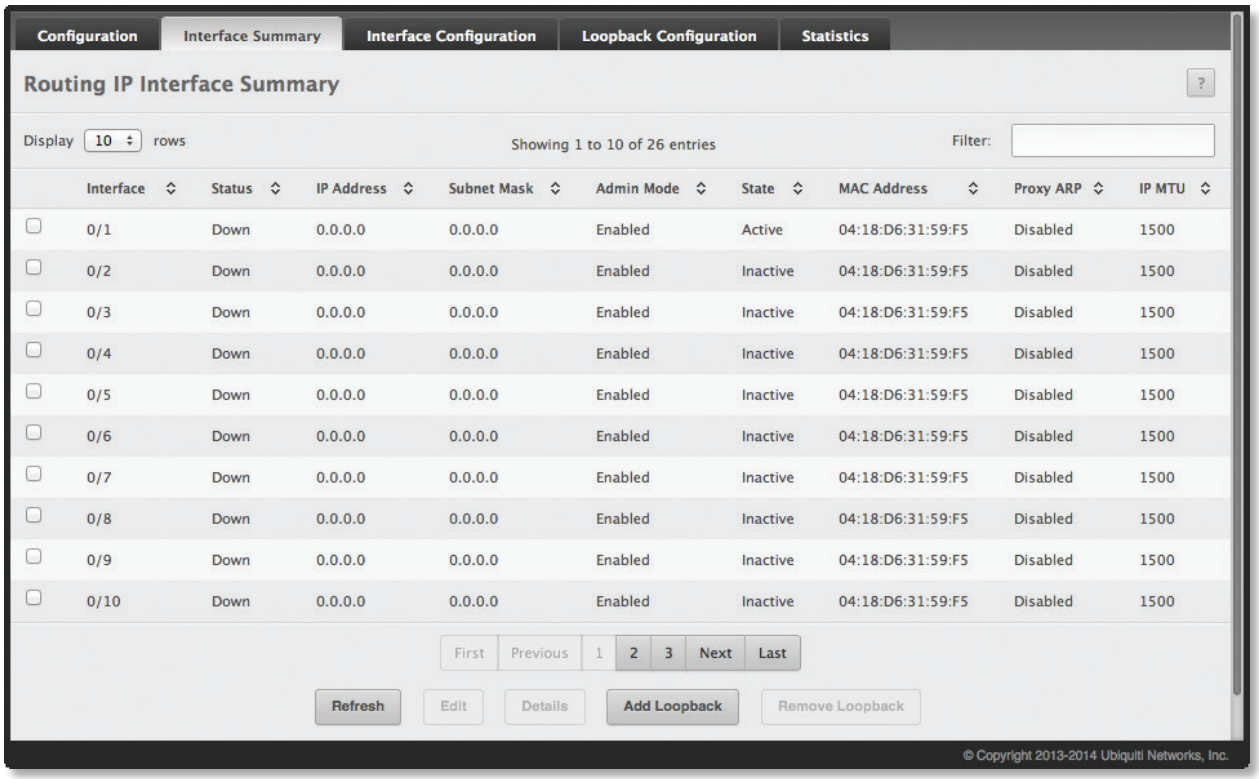
- If you make any changes to the page, click **Submit** to apply the changes to the system.
- Click **Refresh** to refresh the page with the most current data from the switch.

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.



Routing IP Interface Summary

Routing IP Interface Summary Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When viewing details about the routing settings for an interface, this field identifies the interface being viewed.
Status	Indicates whether the interface is capable of routing IP packets (*Up*) or cannot route packets (*Down*). For the status to be Up, the routing mode and administrative mode for the interface must be enabled. Additionally, the interface must have an IP address and be physically up (active link).
IP Address	The IP address of the interface.
Subnet Mask	The IP subnet mask for the interface (also known as the network mask or netmask). It defines the portion of the interface's IP address that is used to identify the attached network.
Admin Mode	The administrative mode of the interface, which is either *Enabled* or *Disabled*.
State	The state of the interface, which is either *Active* or *Inactive*. An interface is considered active if the link is up, and the interface is in a forwarding state.
MAC Address	The burned-in physical address of the interface. The format is six two-digit hexadecimal numbers separated by colons, for example 00:06:29:32:81:40.
Proxy ARP	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.
IP MTU	The largest IP packet size the interface can transmit, in bytes. The IP Maximum Transmission Unit (MTU) is the maximum frame size minus the length of the Layer-2 header.

Routing IP Interface Summary Fields (Continued)

Field	Description
Details window – If you select an interface and click **Details**, the *Details* window opens and displays the following additional routing information for the selected interface:	
Routing Mode	Indicates whether routing is administratively *Enabled* or *Disabled* on the interface.
Link Speed Data Rate	The physical link data rate of the interface.
IP Address Configuration Method	The source of the IP address, which is one of the following: • **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.
Bandwidth	The configured bandwidth on this interface. This setting communicates the speed of the interface to higher-level protocols.
Encapsulation Type	The link layer encapsulation type for packets transmitted from the interface, which can be either *Ethernet* or *SNAP*.
Forward Net Directed Broadcasts	Indicates how the interface handles network-directed broadcast packets. A network-directed broadcast is a broadcast directed to a specific subnet. The possible values are as follows: • **Enabled** Network directed broadcasts are forwarded. • **Disabled** Network directed broadcasts are dropped.
Local Proxy ARP	Indicates whether local proxy ARP is *Enabled* or *Disabled* on the interface. When local proxy ARP is enabled, the interface can respond to an ARP request for a host other than itself. Unlike proxy ARP, local proxy ARP allows the interface to respond to ARP requests for a host that is on the same subnet as the host that sent the ARP request. This feature is useful when a host is not permitted to reply to an ARP request from another host in the same subnet, for example when using the protected ports feature.
Destination Unreachables	Displays *Enabled* 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 *Disabled*, this interface will not send ICMP Destination Unreachable messages to inform the host about the error in reaching the intended destination.
ICMP Redirects	Displays *Enabled* if the interface is allowed to send ICMP Redirect messages; otherwise, displays *Disabled*. 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:

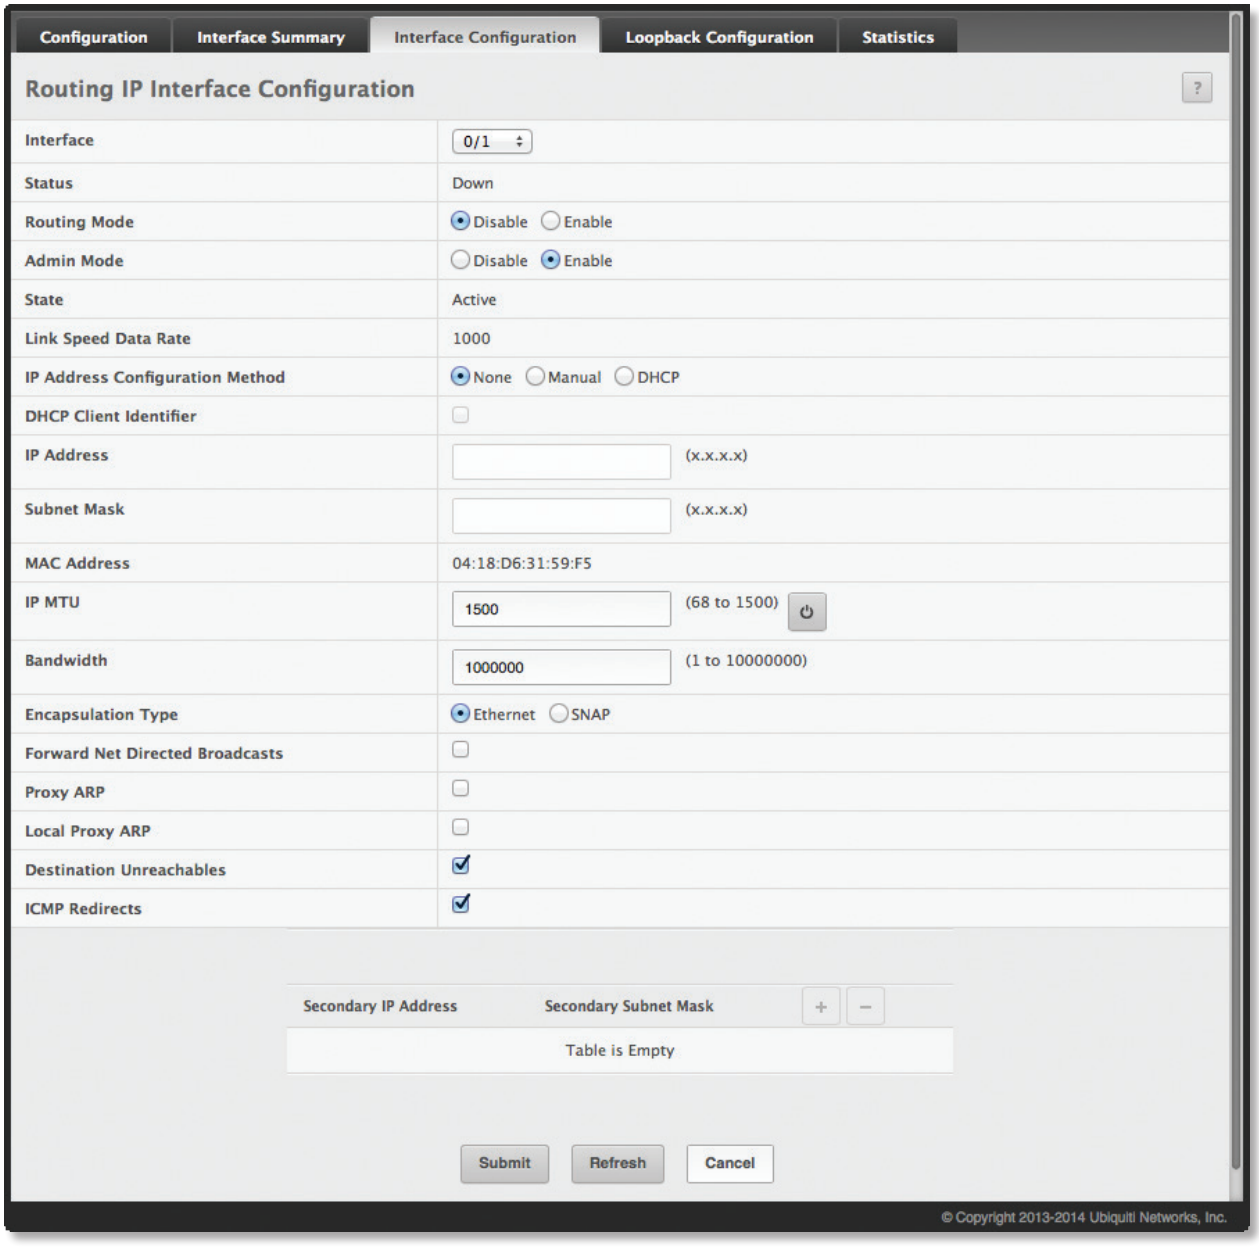
- To edit an interface's routing configuration, select the interface and click **Edit**. The display changes to the *Routing IP Interface Configuration* page; for instructions on using this page, see **"Routing IP Interface Configuration" on page 195**.
- To view detailed routing information on an interface, 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
Interface	The menu contains all interfaces that can be configured for routing. To configure routing settings for an interface, select it from the menu and then configure the rest of the settings on the page.
Status	Indicates whether the interface is currently capable of routing IP packets (*Up*) or cannot route packets (*Down*). For the status to be *Up*, the routing mode and administrative mode for the interface must be enabled. Additionally, the interface must have an IP address and be physically up (active link).
Routing Mode	Used to *Enable* or *Disable* the administrative mode of IP routing on the interface.

Routing IP Interface Configuration Fields (Continued)

Field	Description
Admin Mode	The administrative mode of the interface. If set to *Disable*, the interface cannot forward traffic.
State	The state of the interface, which is either *Active* or *Inactive*. An interface is considered active if the link is up, and the interface is in a forwarding state.
Link Speed Data Rate	The physical link data rate of the interface.
IP Address Configuration Method	The method to use for configuring an IP address on the interface, which can be one of the following: • **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.
IP Address	The IP address of the interface. This field can be configured only when the selected *IP Address Configuration Method* is *Manual*. If the method is *DHCP*, the interface attempts to lease an IP address from a DHCP server on the network, and the IP address appears in this field (read-only) after it is acquired. If this field is blank, the *IP Address Configuration Method* might be *None*, or the method might be *DHCP* and the interface is unable to lease an address.
Subnet Mask	The IP subnet mask for the interface (also known as the network mask or netmask). This field can be configured only when the selected *IP Address Configuration Method* is *Manual*.
MAC Address	The burned-in physical address of the interface. The format is six two-digit hexadecimal numbers separated by colons, for example 00:06:29:32:81:40.
IP MTU	The largest IP packet size the interface can transmit, in bytes. The IP Maximum Transmission Unit (MTU) is the maximum frame size minus the length of the Layer-2 header. Click ⏻ to reset this field.
Bandwidth	The configured bandwidth on this interface. This setting communicates the speed of the interface to higher-level protocols.
Encapsulation Type	The link layer encapsulation type for packets transmitted from the interface: *Ethernet* or *SNAP*.
Forward Net Directed Broadcasts	Determines how the interface handles network-directed broadcast packets. A network-directed broadcast is a broadcast directed to a specific subnet. If this option is selected, network directed broadcasts are forwarded. If this option is cleared, network directed broadcasts are dropped.
Proxy ARP	When this option is selected, proxy ARP is enabled, and the interface can respond to an ARP request for a host other than itself. An interface can act as an ARP proxy if it is aware of the destination and can route packets to the intended host, which is on a different subnet than the host that sent the ARP request.
Local Proxy ARP	When this option is selected, local proxy ARP is enabled, and the interface can respond to an ARP request for a host other than itself. Unlike proxy ARP, local proxy ARP allows the interface to respond to ARP requests for a host that is on the same subnet as the host that sent the ARP request. This feature is useful when a host is not permitted to reply to an ARP request from another host in the same subnet, for example when using the protected ports feature.
Destination Unreachables	When this option is selected, the interface is allowed to send 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.
ICMP Redirects	When this option is selected, the interface is allowed to send ICMP Redirect messages. The device sends an ICMP Redirect message on an interface only if ICMP Redirects are enabled both globally and on the interface. An ICMP Redirect message notifies a host when a better route to a particular destination is available on the network segment.
Secondary IP Address	To add a secondary IP address on the interface, click + in the header row and enter the address in the appropriate field in the *Secondary IP Address Configuration* window. You can add one or more secondary IP addresses to an interface only if the interface already has a primary IP address. To remove a configured secondary IP address, click the entry's associated − button. To remove *all* configured secondary IP addresses, click − in the header row.
Secondary Subnet Mask	The subnet mask associated with the secondary IP address. You configure this field in the *Secondary IP Address Configuration* 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.



Routing IP Statistics

Routing IP Statistics Fields

Field	Description
IpInReceives	The total number of input datagrams received from interfaces, including those received in error.
IpInHdrErrors	The number of input datagrams discarded due to errors in their IP headers, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, errors discovered in processing their IP options, etc.
IpInAddrErrors	The number of input datagrams discarded because the IP address in their IP header's destination field was not a valid address to be received at this entity. This count includes invalid addresses (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.
IpForwDatagrams	The number of input datagrams for which this entity was not their final IP destination, as a result of which an attempt was made to find a route to forward them to that final destination. In entities which do not act as IP Gateways, this counter includes only those packets which were Source-Routed via this entity, and the Source-Route option processing was successful.
IpInUnknownProtos	The number of locally-addressed datagrams received successfully but discarded because of an unknown or unsupported protocol.
IpInDiscards	The number of input IP datagrams for which no problems were encountered to prevent their continued processing, but which were discarded (e.g., for lack of buffer space). Note that this counter does not include any datagrams discarded while awaiting re-assembly.
IpInDelivers	The total number of input datagrams successfully delivered to IP user-protocols (including ICMP).
IpOutRequests	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 *IpForwDatagrams*.
IpOutDiscards	The number of output IP datagrams for which no problem was encountered to prevent their transmission to their destination, but which were discarded (e.g., for lack of buffer space). Note that this counter would include datagrams counted in *IpForwDatagrams* if any such packets met this (discretionary) discard criterion.
IpOutNoRoutes	The number of IP datagrams discarded because no route could be found to transmit them to their destination. Note that this counter includes any packets counted in *IpForwDatagrams* which meet this `no-route' criterion. Note that this includes any datagrams which a host cannot route because all of its default gateways are down.
IpReasmTimeout	The maximum number of seconds which received fragments are held while they are awaiting reassembly at this entity.
IpReasmReqds	The number of IP fragments received which needed to be reassembled at this entity.
IpReasmOKs	The number of IP datagrams successfully re-assembled.
IpReasmFails	The number of failures detected by the IP re-assembly algorithm (for whatever reason: timed out, errors, etc.). Note that this is not necessarily a count of discarded IP fragments since some algorithms can lose track of the number of fragments by combining them as they are received.
IpFragOKs	The number of IP datagrams that have been successfully fragmented at this entity.
IpFragFails	The number of IP datagrams that have been discarded because they needed to be fragmented at this entity but could not be, e.g., because their Don't Fragment flag was set.
IpFragCreates	The number of IP datagram fragments that have been generated as a result of fragmentation at this entity.
IpRoutingDiscards	The number of routing entries which were chosen to be discarded even though they are valid. One possible reason for discarding such an entry could be to free-up buffer space for other routing entries.
IcmpInMsgs	The total number of ICMP messages which the entity received. Note that this counter includes all those counted by IcmpInErrors.
IcmpInErrors	The number of ICMP messages which the entity received but determined as having ICMP-specific errors (bad ICMP checksums, bad length, etc.).
IcmpInDestUnreachs	The number of ICMP Destination Unreachable messages received.
IcmpInTimeExcds	The number of ICMP Time Exceeded messages received.
IcmpInParmProbs	The number of ICMP Parameter Problem messages received.
IcmpInSrcQuenchs	The number of ICMP Source Quench messages received.
IcmpInRedirects	The number of ICMP Redirect messages received.
IcmpInEchos	The number of ICMP Echo (request) messages received.

Routing IP Statistics Fields (Continued)

Field	Description
IcmpInEchoReps	The number of ICMP Echo Reply messages received.
IcmpInTimestamps	The number of ICMP Timestamp (request) messages received.
IcmpInTimestampReps	The number of ICMP Timestamp Reply messages received.
IcmpInAddrMasks	The number of ICMP Address Mask Request messages received.
IcmpInAddrMaskReps	The number of ICMP Address Mask Reply messages received.
IcmpOutMsgs	The total number of ICMP messages which this entity attempted to send. Note that this counter includes all those counted by IcmpOutErrors.
IcmpOutErrors	The number of ICMP messages which this entity did not send due to problems discovered within ICMP such as a lack of buffers. This value should not include errors discovered outside the ICMP layer such as the inability of IP to route the resultant datagram. In some implementations there may be no types of error which contribute to this counter's value.
IcmpOutDestUnreachs	The number of ICMP Destination Unreachable messages sent.
IcmpOutTimeExcds	The number of ICMP Time Exceeded messages sent.
IcmpOutParmProbs	The number of ICMP Parameter Problem messages sent.
IcmpOutSrcQuenchs	The number of ICMP Source Quench messages sent.
IcmpOutRedirects	The number of ICMP Redirect messages sent. For a host, this object is always zero, since hosts do not send redirects.
IcmpOutEchos	The number of ICMP Echo (request) messages sent.
IcmpOutEchoReps	The number of ICMP Echo Reply messages sent.
IcmpOutTimestamps	The number of ICMP Timestamp (request) messages.
IcmpOutTimestampReps	The number of ICMP Timestamp Reply messages sent.
IcmpOutAddrMasks	The number of ICMP Address Mask Request messages sent.

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

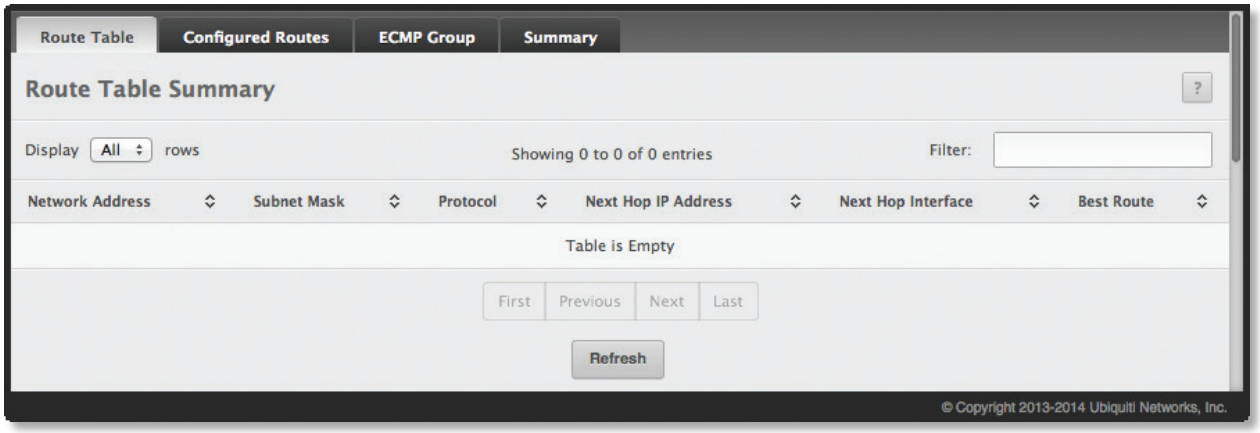
Router

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

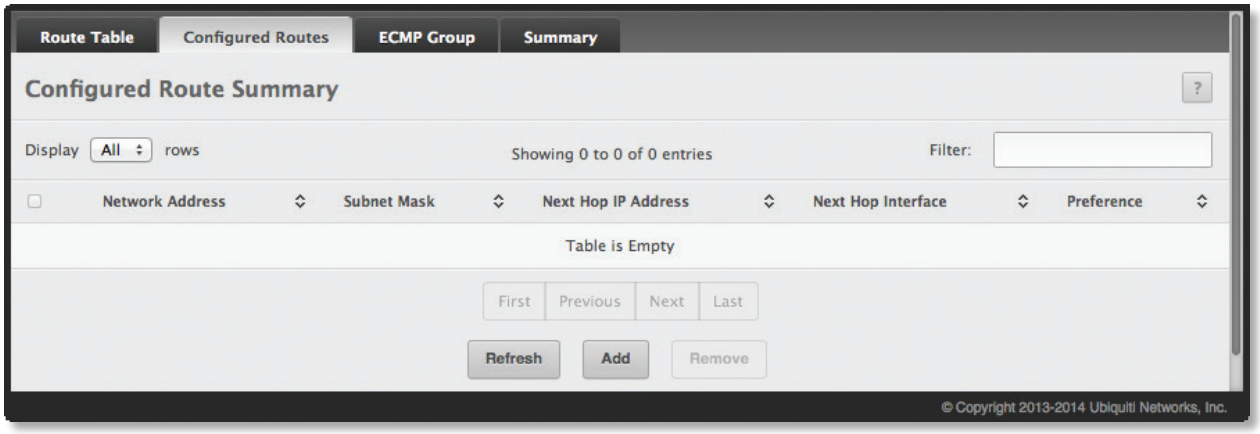
Field	Description
Network Address	The IP route prefix for the destination.
Subnet Mask	Also referred to as the subnet/network mask, this indicates the portion of the IP interface address that identifies the attached network.
Protocol	This field tells which protocol created the specified route. A route 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
Next Hop IP Address	The outgoing router IP address to use when forwarding traffic to the next router (if any) in the path towards the destination. The next router is always one of the adjacent neighbors or the IP address of the local interface for a directly-attached network.
Next Hop Interface	The outgoing interface to use when forwarding traffic to the destination. For a static reject route, the next hop is Null.
Best Route	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
Network Address	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.
Subnet Mask	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.
Next Hop IP Address	The outgoing router IP address to use when forwarding traffic to the next router (if any) in the path towards the destination. The next router is always one of the adjacent neighbors or the IP address of the local interface for a directly-attached network. When adding a static reject route, this field is not available because the packets are dropped rather than forwarded.
Next Hop Interface	The outgoing interface to use when forwarding traffic to the destination. For a static reject route, the next hop is Null.
Preference	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.

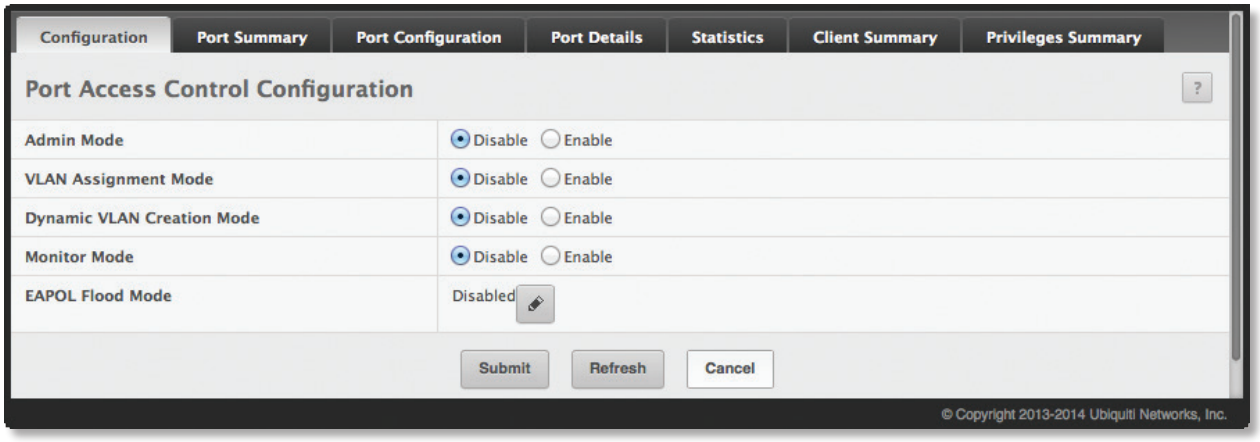
Authentication Server: Specifies the external server, for example, the RADIUS server that performs the authentication on behalf of the authenticator, and indicates whether the user is authorized to access system services.

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
Admin Mode	Specifies whether to *Enable* or *Disable* port-based authentication on the switch. The default is *Disable*.
VLAN Assignment Mode	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.
Dynamic VLAN Creation Mode	The administrative mode of dynamic VLAN creation on the device. Select *Enable* 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.
Monitor Mode	The administrative mode of the Monitor Mode feature on the device. Monitor mode is a special mode that can be enabled in conjunction with port-based access control. Monitor mode provides a way for network administrators to identify possible issues with the port-based access control configuration on the device without affecting the network access to the users of the device. It allows network access even in cases where there is a failure to authenticate, but it logs the results of the authentication process for diagnostic purposes. If the device fails to authenticate a client for any reason (for example, RADIUS access reject from the RADIUS server, RADIUS timeout, or the client itself is 802.1X unaware), the client is authenticated and is undisturbed by the failure condition(s). The reasons for failure are logged and buffered into the local logging database for tracking purposes.
EAPOL Flood Mode	The administrative mode of the Extensible Authentication Protocol (EAP) over LAN (EAPOL) flood support on the device. EAPOL Flood Mode can be enabled when *Admin Mode* and *Monitor Mode* are disabled.

Use the buttons to perform the following tasks:

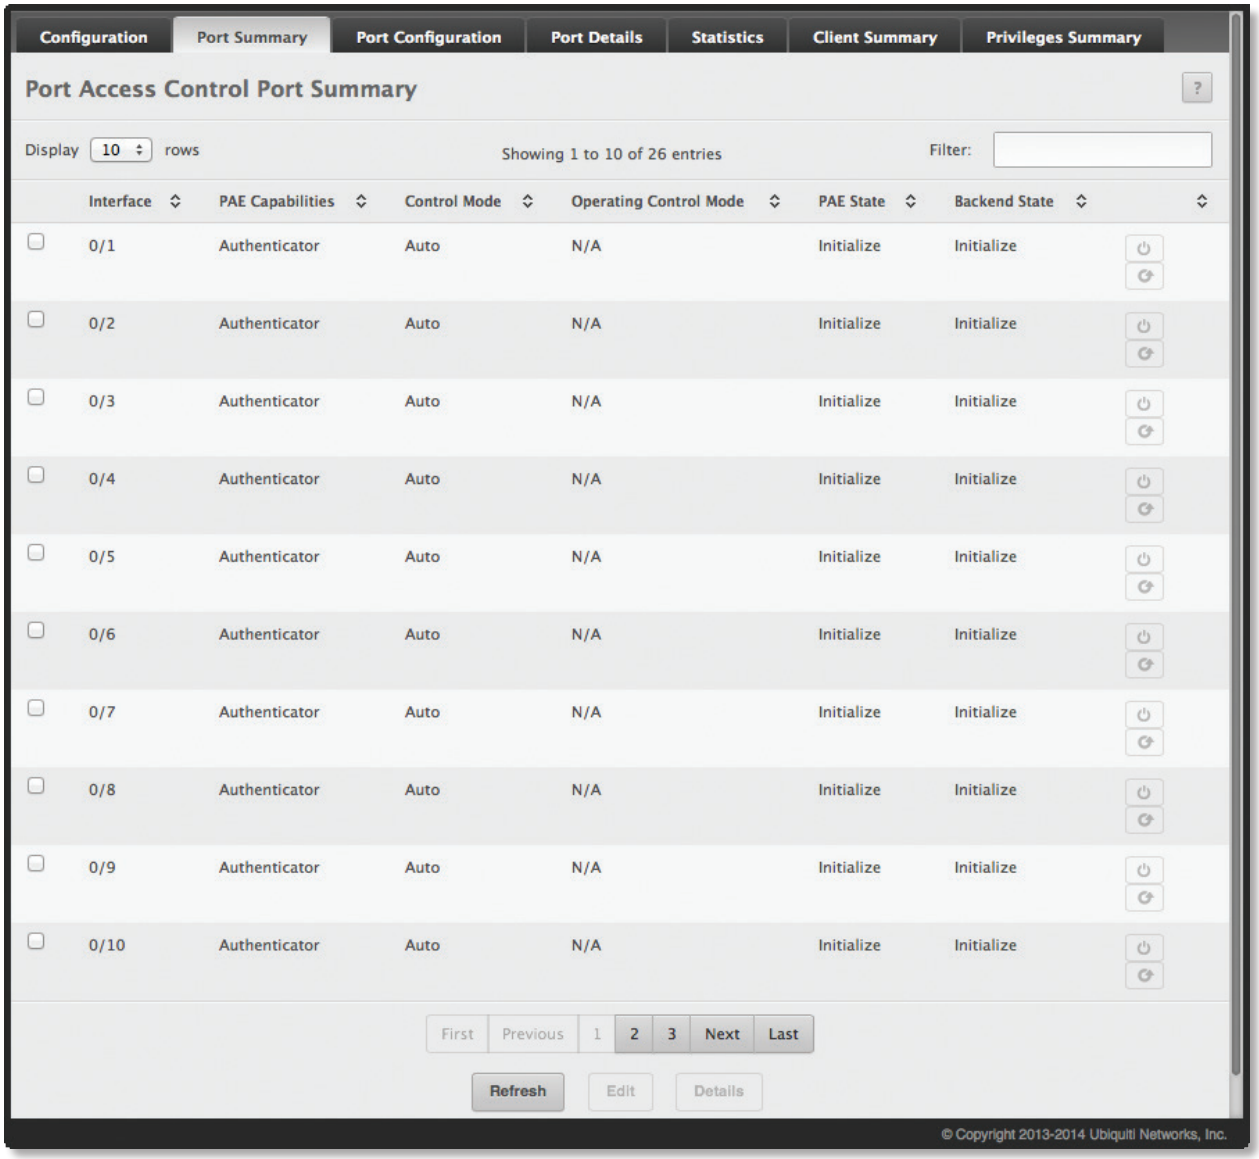
- If you change any settings, click **Submit** to apply the new settings to the system.
- Click **Refresh** to refresh the page with the most current data from the switch.

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.



Port Access Control Port Summary

Port Access Control Port Summary Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.
PAE Capabilities	The Port Access Entity (PAE) role, which is one of the following: • **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
Control Mode	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.
Operating Control Mode	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.
PAE State	The current state of the authenticator PAE state machine, which is the 802.1X process that controls access to the port. The state can be one of the following: • **Initialize** • **Disconnected** • **Connecting** • **Authenticating** • **Authenticated** • **Aborting** • **Held** • **ForceAuthorized** • **ForceUnauthorized**
Backend State	The current state of the back-end authentication state machine, which is the 802.1X process that controls the interaction between the 802.1X client on the local system and the remote authentication server. The state can be one of the following: • **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 *Control Mode* is **Auto**. Click this button to force the associated interface to restart the authentication process.

Use the buttons to perform the following tasks:

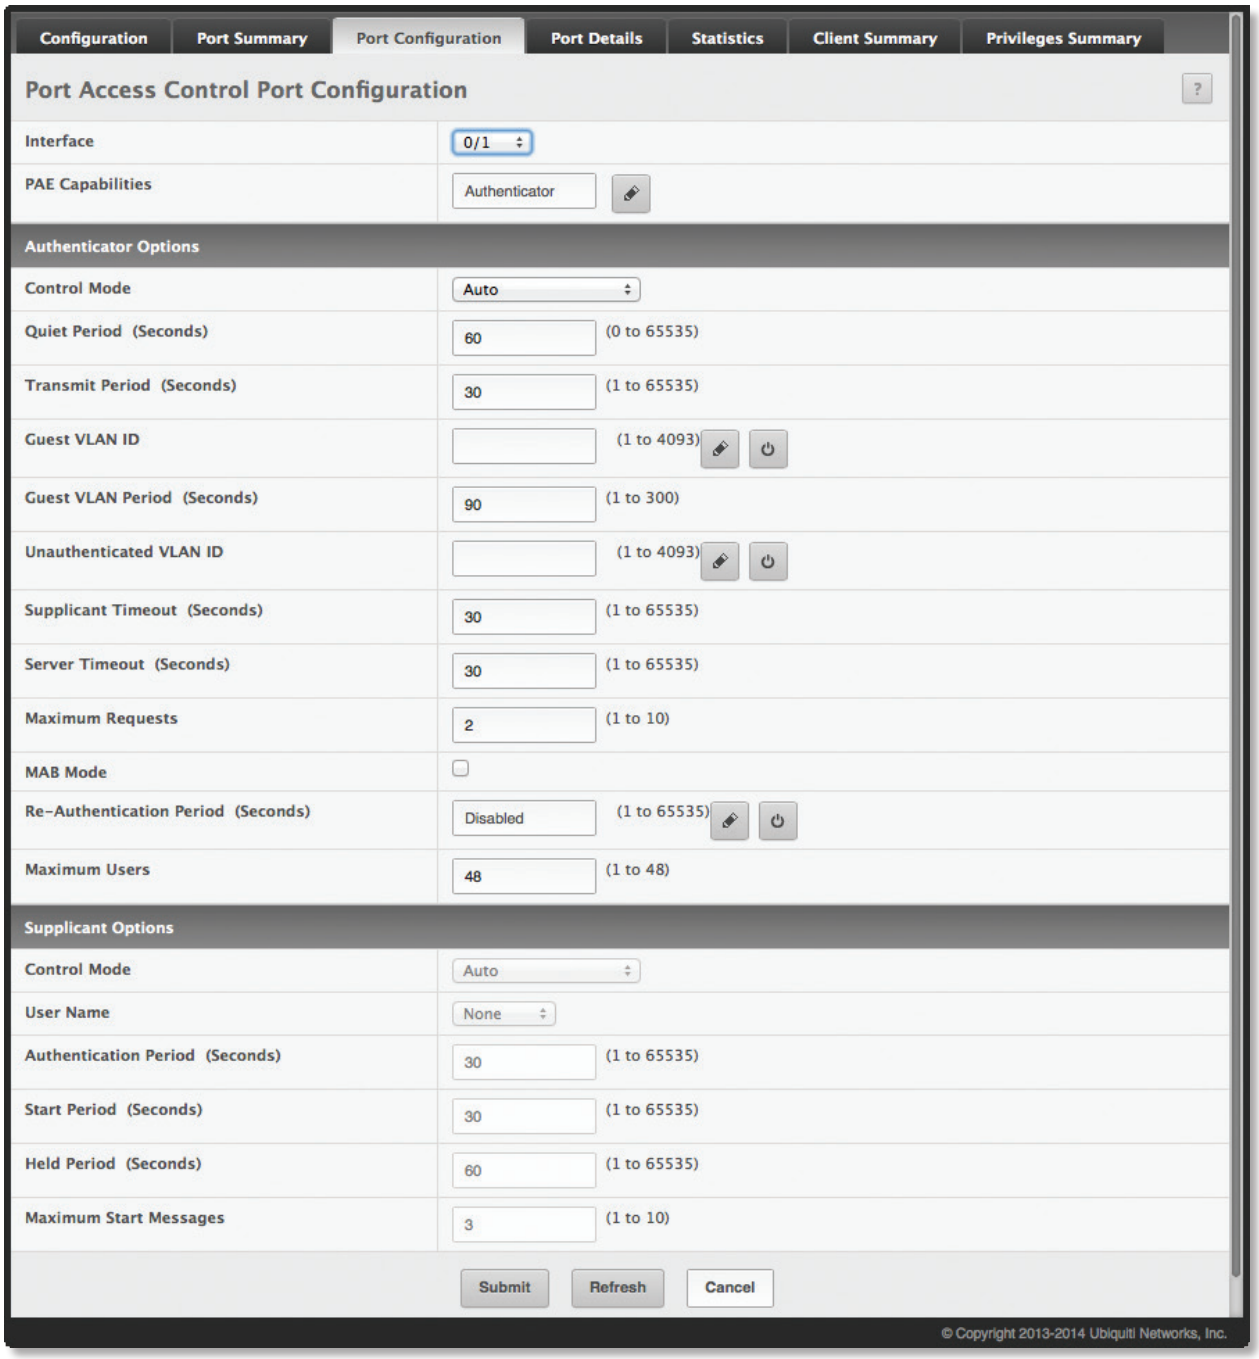
- To change the port-based access control settings for a port, select the port and click **Edit**. The UI automatically redirects to the *Port Access Control Port Configuration* page for the selected port. For information on using this page, refer to **"Global Port Access Control Configuration" on page 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.



Port Access Control Port Configuration

Port Access Control Port Configuration Fields

Field	Description
Interface	The interface with the settings to view or configure. If you have been redirected to this page, this field is read-only and displays the interface that was selected on the *Port Access Control Port Summary* page.
PAE Capabilities	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 *Set PAE Capabilities* window.
Authenticator Options – The fields in this section can be changed only when the selected port is configured as an authenticator port (that is, the *PAE Capabilities* field is set to **Authenticator**).	
Control Mode	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.
Quiet Period	The number of seconds that the port remains in the quiet state following a failed authentication exchange.
Transmit Period	The value, in seconds, of the timer used by the authenticator state machine on the port to determine when to send an EAPOL EAP Request/Identity frame to the supplicant.
Guest VLAN ID	The VLAN ID 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 *Guest VLAN ID*. Click this button to reset the *Guest VLAN ID* to the default value.
Guest VLAN Period	The value, in seconds, of the timer used for guest VLAN authentication.
Unauthenticated VLAN ID	The VLAN ID of the unauthenticated VLAN. Hosts that fail the authentication might be denied access to the network or placed on a VLAN created for unauthenticated clients. This VLAN might be configured with limited network access. Click this button to set the *Unauthenticated VLAN ID*. Click this button to reset the *Unauthenticated VLAN ID* to the default value.
Supplicant Timeout	The amount of time that the port waits for a response before retransmitting an EAP request frame to the client.
Server Timeout	The amount of time the port waits for a response from the authentication server.
Maximum Requests	The maximum number of times that the port sends an EAP request frame (assuming that no response is received) to the client before restarting the authentication process.
MAB Mode	The MAC-based Authentication Bypass (MAB) mode on the port, which can be enabled or disabled.
Re-Authentication Period	The amount of time that clients can be connected to the port without being reauthenticated. If this field is disabled, connected clients are not forced to reauthenticate periodically. Click this button to set the *Re-Authentication Period*. Click this button to reset the *Re-Authentication Period* to the default value.
Maximum Users	The maximum number of clients supported on the port if the *Control Mode* on the port is **MAC-Based** 802.1X authentication.

Port Access Control Port Configuration Fields (Continued)

Field	Description
Supplicant Options – The fields in this section can be changed only when the selected port is configured as a supplicant port (that is, the *PAE Capabilities* field is set to **Supplicant**).	
Control Mode	The port-based access control mode on the port, which is one of the following: • **Auto** The port is in an unauthorized state until a successful authentication exchange has taken place between the supplicant port, the authenticator port, and the authentication server. • **Force Unauthorized** The port is placed into an unauthorized state and is automatically denied system access. • **Force Authorized** The port is placed into an authorized state and does not require client port-based authentication to be able to send and receive traffic.
User Name	The name the port uses to identify itself as a supplicant to the authenticator port. The menu includes the users that are configured for system management. When authenticating, the supplicant provides the password associated with the selected *User Name*.
Authentication Period	The amount of time the supplicant port waits to receive a challenge from the authentication server. If the configured *Authentication Period* expires, the supplicant retransmits the authentication request until it is authenticated or has sent the number of messages configured in the *Maximum Start Messages* field.
Start Period	The amount of time the supplicant port waits for a response from the authenticator port after sending a Start packet. If no response is received, the supplicant retransmits the Start packet.
Held Period	The amount of time the supplicant port waits before contacting the authenticator port after an active 802.1X session fails.
Maximum Start Messages	The maximum number of Start packets the supplicant port sends to the authenticator port without receiving a response before it considers the authenticator to be 802.1X-unaware.

Use the buttons to perform the following tasks:

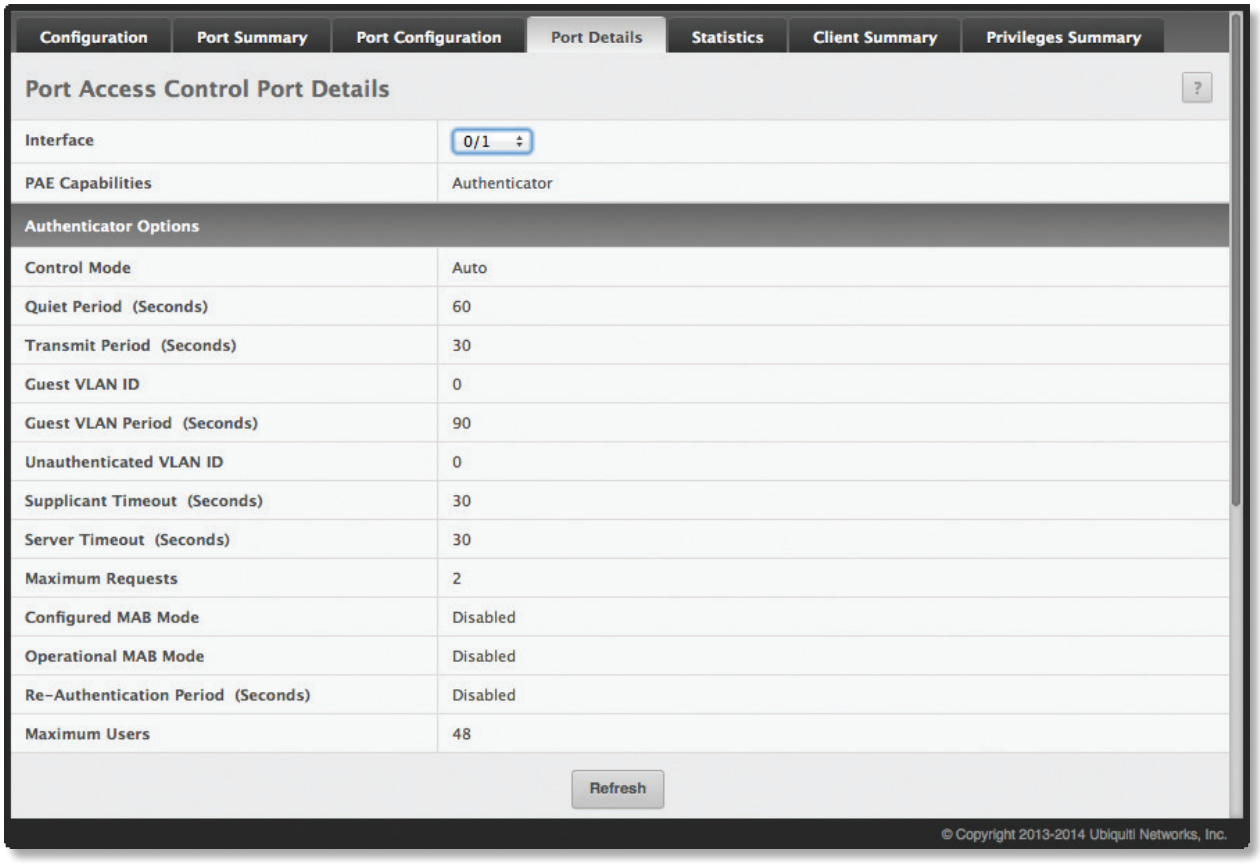
- If you change any settings on this page, click **Submit** to apply the changes.
- Click **Refresh** to refresh the page with the most current data from the switch.

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.



Port Access Control Port Details

Port Access Control Port Details Fields

Field	Description
Interface	The interface associated with the rest of the data on the page.
PAE Capabilities	The Port Access Entity (PAE) role, which is one of the following: • **Authenticator** The port enforces authentication and passes authentication information from a remote supplicant (client or host) to the authentication server. If the server successfully authenticates the supplicant, the port allows access. • **Supplicant** The port is connected to an authenticator port and must be granted permission by the authentication server before it can send and receive traffic through the remote port.
Authenticator Options – The fields in this section provide information about the settings that apply to the port when it is configured as an 802.1X authenticator.	
Control Mode	The port-based access control mode on the port, which is one of the following: • **Auto** The port is unauthorized until a successful authentication exchange has taken place. • **Force Unauthorized** The port ignores supplicant authentication attempts and does not provide authentication services to the client. • **Force Authorized** The port sends and receives normal traffic without client port-based authentication. • **MAC-Based** This mode allows multiple supplicants connected to the same port to each authenticate individually. Each host connected to the port must authenticate separately in order to gain access to the network. The hosts are distinguished by their MAC addresses.
Quiet Period	The number of seconds that the port remains in the quiet state following a failed authentication exchange.

Port Access Control Port Details Fields (Continued)

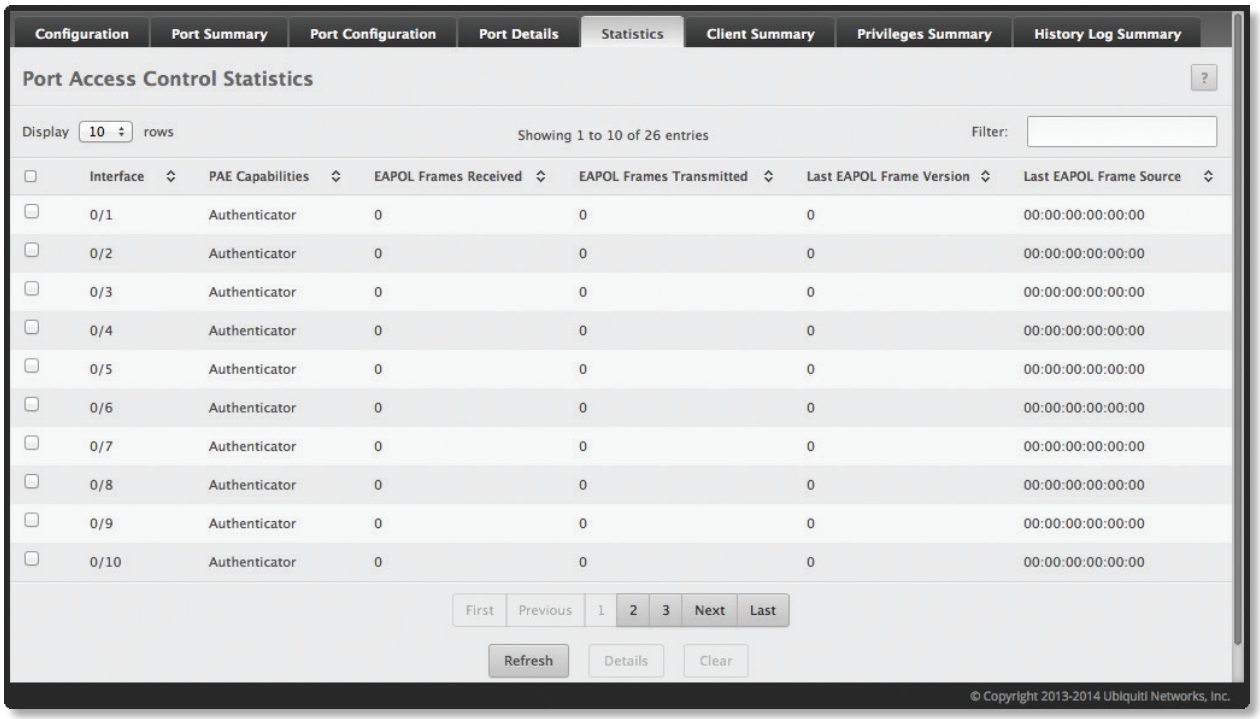
Field	Description
Transmit Period	The value, in seconds, of the timer used by the authenticator state machine on the port to determine when to send an EAPOL EAP Request/Identity frame to the supplicant.
Guest VLAN ID	The VLAN ID for the guest VLAN. The guest VLAN allows the port to provide a distinguished service to unauthenticated users. This feature provides a mechanism to allow users access to hosts on the guest VLAN.
Guest VLAN Period	The value, in seconds, of the timer used for guest VLAN authentication.
Unauthenticated VLAN ID	The VLAN ID of the unauthenticated VLAN. Hosts that fail the authentication might be denied access to the network or placed on a VLAN created for unauthenticated clients. This VLAN might be configured with limited network access.
Supplicant Timeout	The amount of time that the port waits for a response before retransmitting an EAP request frame to the client.
Server Timeout	The amount of time the port waits for a response from the authentication server.
Maximum Requests	The maximum number of times that the port sends an EAP request frame (assuming that no response is received) to the client before restarting the authentication process.
Configured MAB Mode	*The configured MAC-based Authentication Bypass (MAB) mode on the port.*
Operational MAB Mode	*The operational MAC-based Authentication Bypass (MAB) mode on the port.*
Re-Authentication Period	The amount of time that clients can be connected to the port without being reauthenticated. If this field is disabled, connected clients are not forced to reauthenicate periodically.
Maximum Users	The maximum number of clients supported on the port if the Control Mode on the port is MAC-based 802.1X authentication.

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.



Port Access Control Statistics

Port Access Control Statistics Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When viewing detailed information for an interface, this field identifies the interface being viewed.
PAE Capabilities	The Port Access Entity (PAE) role, which is one of the following: • **Authenticator** The port enforces authentication and passes authentication information from a remote supplicant (similar to a client or host) to the authentication server. If the server successfully authenticates the supplicant, the port allows access. • **Supplicant** The port must be granted permission by the authentication server before it can access the remote authenticator port.
EAPOL Frames Received	The total number of valid EAPOL frames received on the interface.
EAPOL Frames Transmitted	The total number of EAPOL frames sent by the interface.
Last EAPOL Frame Version	The protocol version number attached to the most recently received EAPOL frame.
Last EAPOL Frame Source	The source MAC address attached to the most recently received EAPOL frame.
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 *Details* window. The fields this window displays depend on whether the interface is configured as an authenticator or supplicant, as noted in the applicable field descriptions.	
EAPOL Start Frames Received	The total number of EAPOL-Start frames received on the interface. EAPOL-Start frames are sent by a supplicant to initiate the 802.1X authentication process when it connects to the interface. This field is displayed only if the interface is configured as an authenticator.
EAPOL Logoff Frames Received	The total number of EAPOL-Logoff frames received on the interface. EAPOL-Logoff frames are sent by a supplicant to indicate that it is disconnecting from the network, and the interface can return to the unauthorized state. This field is displayed only if the interface is configured as an authenticator.

Port Access Control Statistics Fields (Continued)

Field	Description
EAP Response/ID Frames Received	The total number of EAP-Response Identity frames the interface has received. EAP-Response Identity frames are sent by a supplicant to provide user information that is used to for authentication. This field is displayed only if the interface is configured as an authenticator.
EAP Response Frames Received	The total number of EAP-Response frames the interface has received. EAP-Response frames are sent from a supplicant to an authentication server during the authentication process. This field is displayed only if the interface is configured as an authenticator.
EAP Request/ID Frames Transmitted	The total number of EAP-Request Identity frames the interface has sent. EAP-Request Identity frames are sent from an authenticator to a supplicant to request user information that is used to for authentication. This field is displayed only if the interface is configured as an authenticator.
EAPOL Start Frames Transmitted	The total number of EAPOL-Start frames the interface has sent to a remote authenticator. EAPOL-Start frames are sent by a supplicant to initiate the 802.1X authentication process when it connects to the interface. This field is displayed only if the interface is configured as a supplicant.
EAPOL Logoff Frames Transmitted	The total number of EAPOL-Logoff frames the interface has sent to a remote authenticator. EAPOL-Logoff frames are sent by a supplicant to indicate that it is disconnecting from the network, and the interface can return to the unauthorized state. This field is displayed only if the interface is configured as a supplicant.
EAP Response/ID Frames Transmitted	The total number of EAP-Response Identity frames the interface has sent. EAP-Response Identity frames are sent by a supplicant to provide user information that is used to for authentication. This field is displayed only if the interface is configured as a supplicant.
EAP Request/ID Frames Received	The total number of EAP-Request Identity frames the interface has received. EAP-Request Identity frames are sent from an authenticator to a supplicant to request user information that is used to for authentication. This field is displayed only if the interface is configured as a supplicant.
EAP Request Frames Received	The total number of EAP-Request frames the interface has received. EAP-Request frames are sent from the authentication server to the supplicant during the authentication process. This field is displayed only if the interface is configured as a supplicant.
Invalid EAPOL Frames Received	The number of unrecognized EAPOL frames received on the interface.
EAPOL Length Error Frames Received	The number of EAPOL frames with an invalid packet body length received on the interface.
Clear (Button)	Resets all statistics counters to 0 for the selected interface or interfaces.

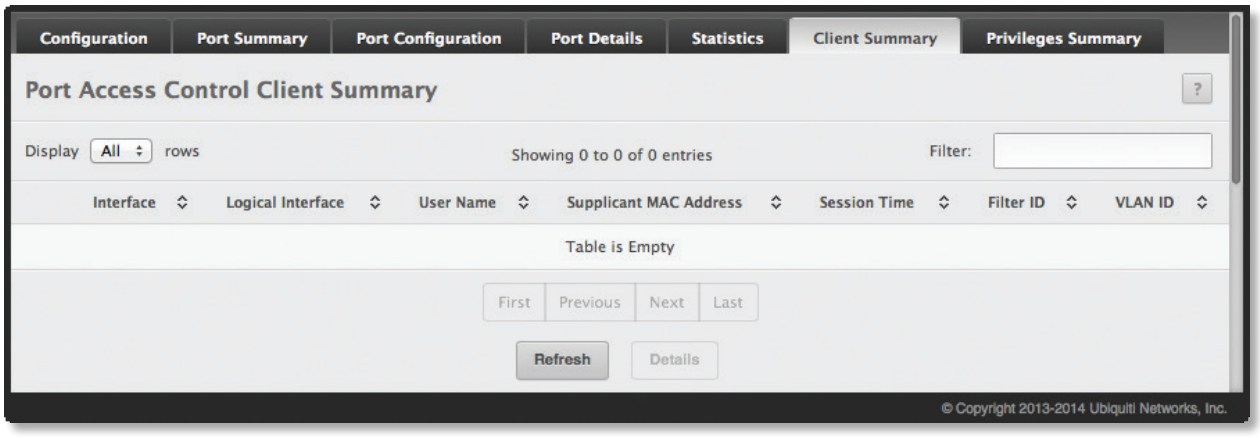
Use the buttons to perform the following tasks:

- 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
Interface	The local interface associated with the rest of the data in the row. When viewing detailed information for an interface, this field identifies the interface being viewed.
Logical Interface	The logical port number associated with the supplicant that is connected to the port.
User Name	The name the client uses to identify itself as a supplicant to the authentication server.
Supplicant MAC Address	The MAC address of the supplicant that is connected to the port.
Session Time	The amount of time that has passed since the connected supplicant was granted access to the network through the authenticator port.
Filter ID	The policy filter ID assigned by the authenticator to the supplicant device.
VLAN ID	The ID of the VLAN the supplicant was placed in as a result of the authentication process.
Details Window Fields After you click **Details**, a window opens and displays the following additional information about the client:	
Session Timeout	The reauthentication timeout period set by the RADIUS server to the supplicant device.
Session Termination Action	The termination action set by the RADIUS server that indicates the action that will take place once the supplicant reaches the session timeout value.

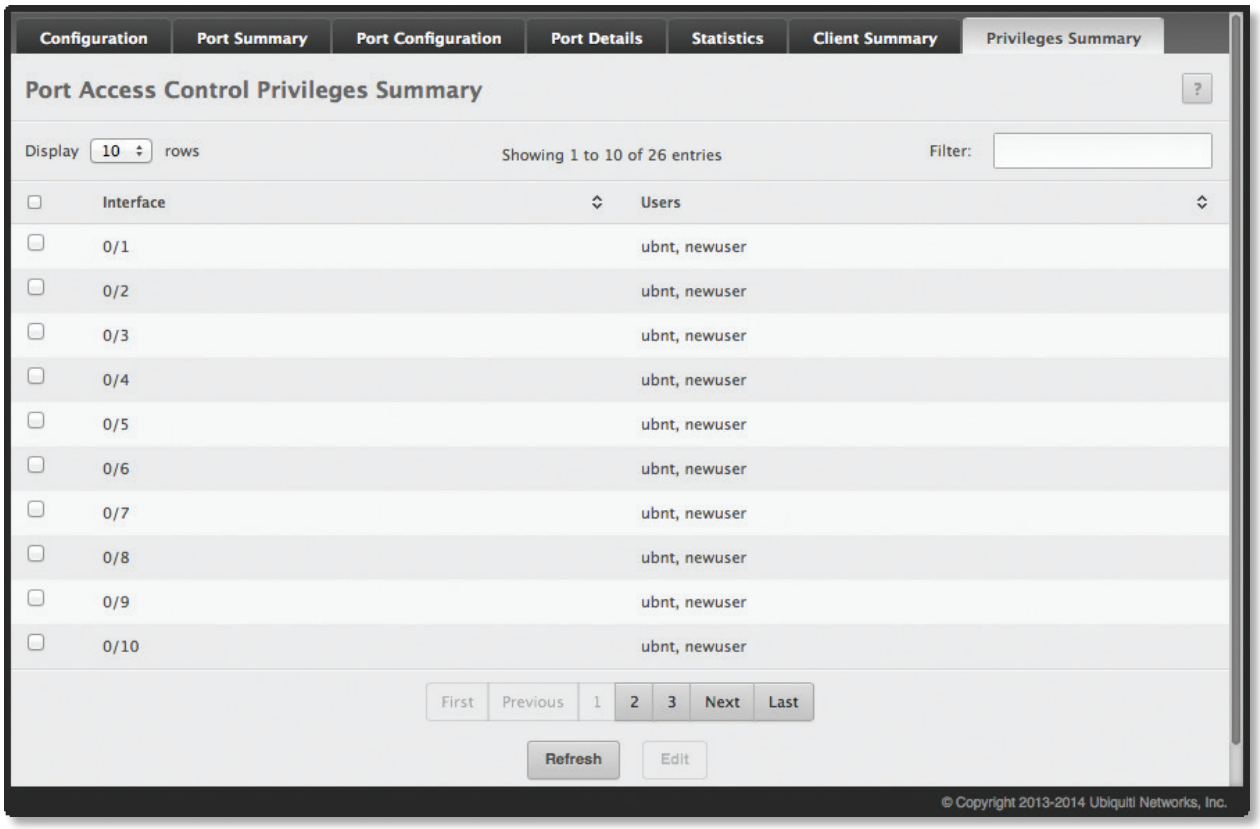
Use the buttons to perform the following tasks:

- 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
Interface	The local interface associated with the rest of the data in the row. When configuring access information for one or more interfaces, this field identifies each interface being configured.
Users	The users that are allowed access to the system through the associated port. When configuring user access for a port, the *Available Users* field lists the users configured on the system who are denied access to the port, while the *Selected Users* 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:

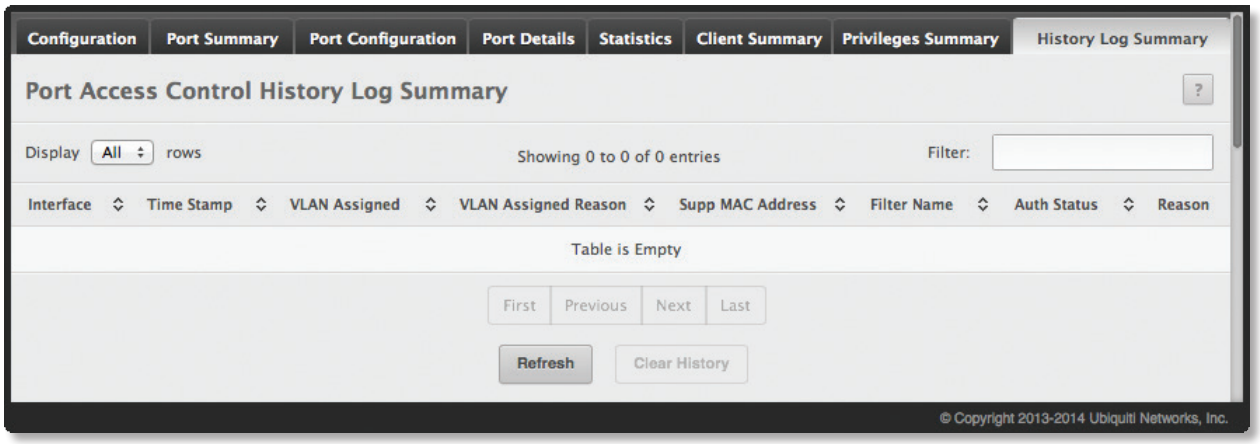
- To change the access control privileges for one or more ports, select the interface(s) to configure and click **Edit**. Configure the settings as needed and click **Submit** to apply the changes. The same settings are applied to all selected interfaces.
- Click **Refresh** to refresh the page with 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
Interface	The interface associated with the rest of the data in the row. Only interfaces that have entries in the log history are listed.
Time Stamp	The absolute time when the authentication event took place.
VLAN Assigned	The ID of the VLAN the supplicant was placed in as a result of the authentication process.
VLAN Assigned Reason	The reason why the authenticator placed the supplicant in the VLAN. Possible values are: • **RADIUS** • **Unauth** • **Default** • **Not Assigned**
Supp MAC Address	The MAC address of the supplicant that is connected to the port.
Filter Name	The policy filter ID assigned by the authenticator to the supplicant device.
Auth Status	The authentication status of the client or port.
Reason	The reason for the successful or unsuccessful authentication.

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

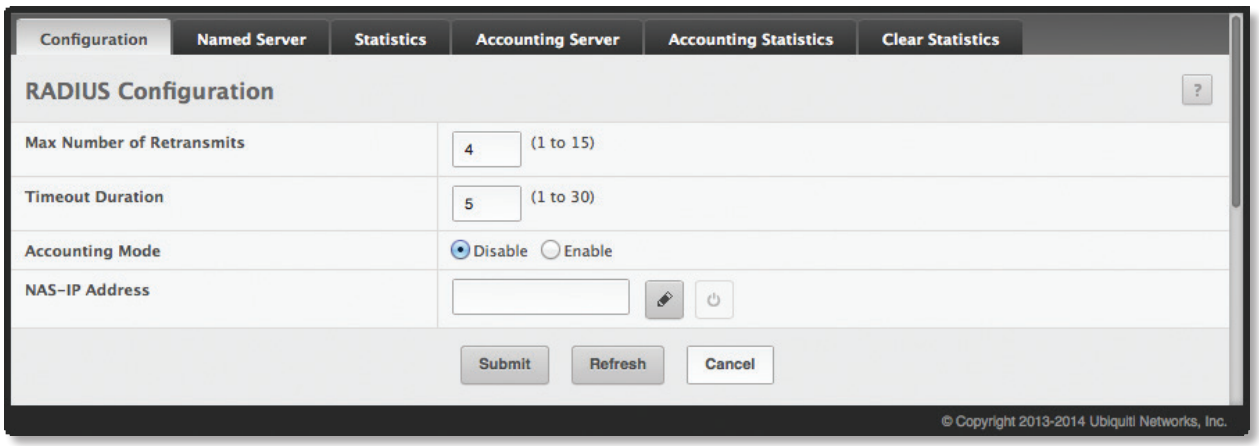
Remote Authorization Dial-In User Service (RADIUS) servers provide additional security for networks. The RADIUS server maintains a user database, which contains per-user authentication information. RADIUS servers provide a centralized authentication method for:

- 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
Max Number of Retransmits	The maximum number of times the RADIUS client on the device will retransmit a request packet to a configured RADIUS server after a response is not received. If multiple RADIUS servers are configured, the max retransmit value will be exhausted on the first server before the next server is attempted. A retransmit will not occur until the configured timeout value on that server has passed without a response from the RADIUS server. Therefore, the maximum delay in receiving a response from the RADIUS server equals the sum of (retransmit × timeout) for all configured servers. If the RADIUS request was generated by a user login attempt, all user interfaces will be blocked until the RADIUS application returns a response.
Timeout Duration	The number of seconds the RADIUS client waits for a response from the RADIUS server. Consideration to maximum delay time should be given when configuring RADIUS timeout and RADIUS max retransmit values.
Accounting Mode	Used to *Enable* or *Disable* the RADIUS accounting mode on the device.
NAS-IP Address	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 *NAS IP Address* to the default value.

Use the buttons at the bottom of the page to perform the following actions:

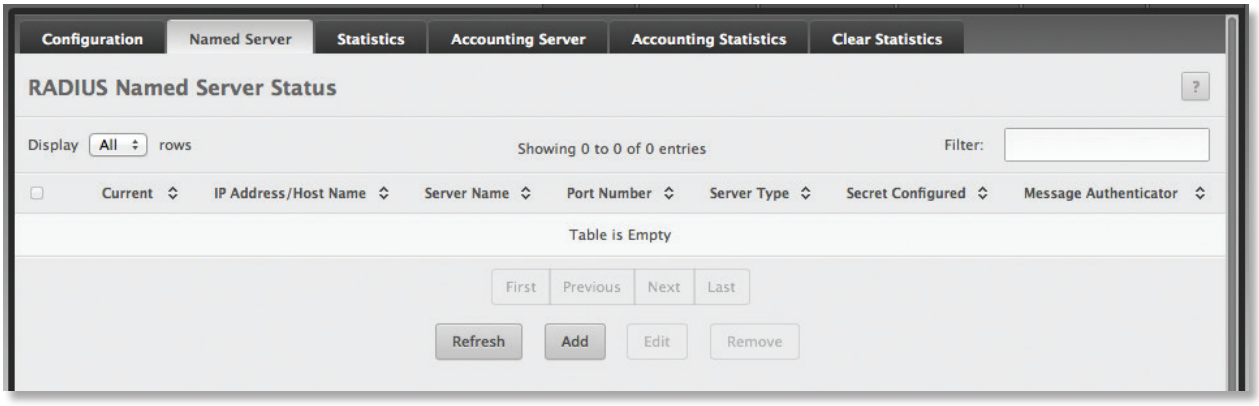
- If you make changes to the page, click **Submit** to apply the changes to the system.
- Click **Refresh** to update the page with the most current information.

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
Current	An asterisk (*) in the column Indicates that the server is the current server for the authentication server group. If no asterisk is present, the server is a backup server. If more than one RADIUS server is configured with the same name, the switch selects one of the servers to be the current server from the group of servers with the same name. When the switch sends a RADIUS request to the named server, the request is directed to the server selected as the current server. Initially the primary server is selected as the current server. If the primary server fails, one of the other servers becomes the current server.
IP Address/Host Address	Shows the IP address of the RADIUS server.
Server Name	Shows the RADIUS server name. Multiple RADIUS servers can have the same name. In this case, RADIUS clients can use RADIUS servers with the same name as backups for each other.
Port Number	Identifies the authentication port the server uses to verify the RADIUS server authentication. The port is a UDP port.
Server Type	Shows whether the server is a *Primary* or *Secondary* server.
Secret Configured	Indicates whether the shared secret for this server has been configured.
Message Authenticator	Shows whether the message authenticator attribute for the selected server is enabled or disabled.

Use the buttons to perform the following tasks:

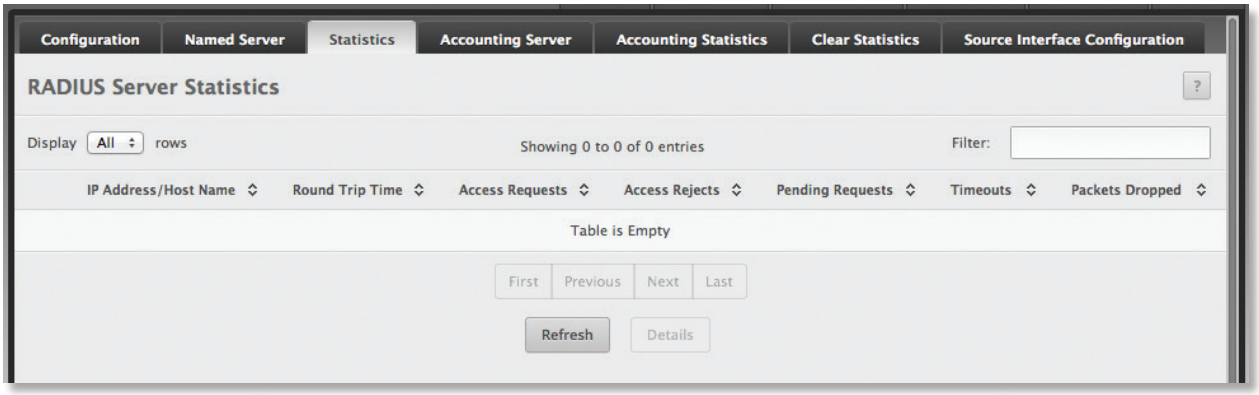
- To add a RADIUS authentication server to the list of servers the RADIUS client can contact, click **Add**, configure the fields, and click **Submit** to apply the changes.
- To change the settings for a configured RADIUS server, select the entry to edit, click **Edit**, configure the fields, and click **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

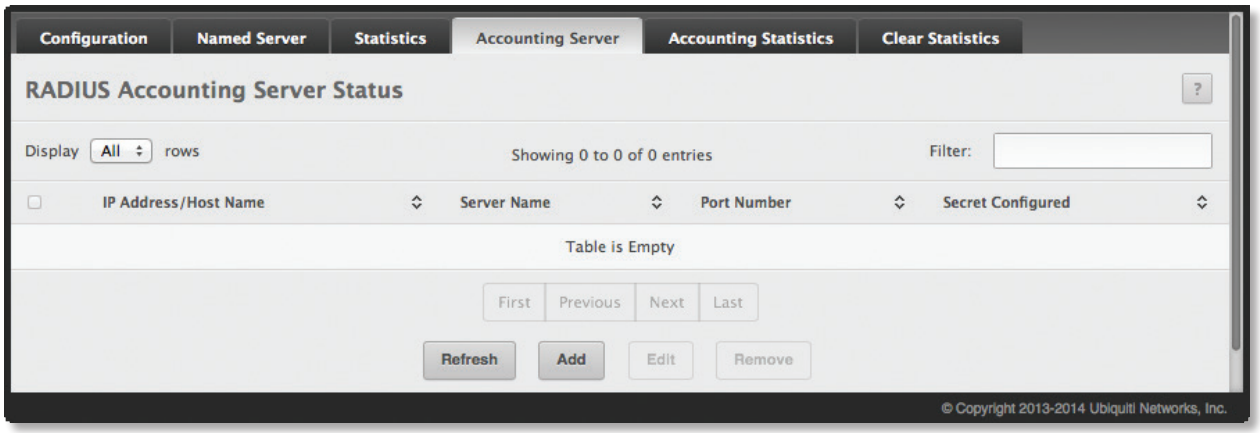
Field	Description
IP Address/Host Name	The IP address or host name of the RADIUS server associated with the rest of the data in the row. When viewing the detailed statistics for a RADIUS server, this field identifies the RADIUS server.
Round Trip Time	The time interval, in hundredths of a second, between the most recent Access-Reply/Access-Challenge and the Access-Request that matched it from the RADIUS authentication server.
Access Requests	The number of RADIUS Access-Request packets sent to the server. This number does not include retransmissions.
Access Rejects	The number of RADIUS Access-Reject packets, including both valid and invalid packets, that were received from the server.
Pending Requests	The number of RADIUS Access-Request packets destined for the server that have not yet timed out or received a response.
Timeouts	The number of times a response was not received from the server within the configured timeout value.
Packets Dropped	The number of RADIUS packets received from the server on the authentication port and dropped for some other reason.
When you click **Details**, the *RADIUS Server Detailed Statistics* 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:	
Access Retransmissions	The number of RADIUS Access-Request packets that had to be retransmitted to the server because the initial Access-Request packet failed to be successfully delivered.
Access Accepts	The number of RADIUS Access-Accept packets, including both valid and invalid packets, that were received from the server.
Access Challenges	The number of RADIUS Access-Challenge packets, including both valid and invalid packets, that were received from the server.
Malformed Access Responses	The number of malformed RADIUS Access-Response packets received from the server. Malformed packets include packets with an invalid length. Bad authenticators, signature attributes, and unknown types are not included as malformed access responses.
Bad Authenticators	The number of RADIUS Access-Response packets containing invalid authenticators or signature attributes received from the server.
Unknown Types	The number of RADIUS packets of unknown type which were received from the server on the authentication port.

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
IP Address/Host Name	The IP address or host name of the RADIUS accounting server. Host names must be resolvable by DNS and are composed of a series of labels separated by dots.
Server Name	The name of the RADIUS accounting server. RADIUS servers that are configured with the same name are members of the same named RADIUS server group. RADIUS accounting servers in the same group serve as backups for each other.
Port Number	The UDP port on the RADIUS accounting server to which the local RADIUS client sends request packets.
Secret Configured	Indicates whether the shared secret for this server has been configured.
Secret	The shared secret text string used for authenticating and encrypting all RADIUS communications between the RADIUS client on the device and the RADIUS accounting server. The secret specified in this field must match the shared secret configured on the RADIUS accounting server.

Use the buttons to perform the following tasks:

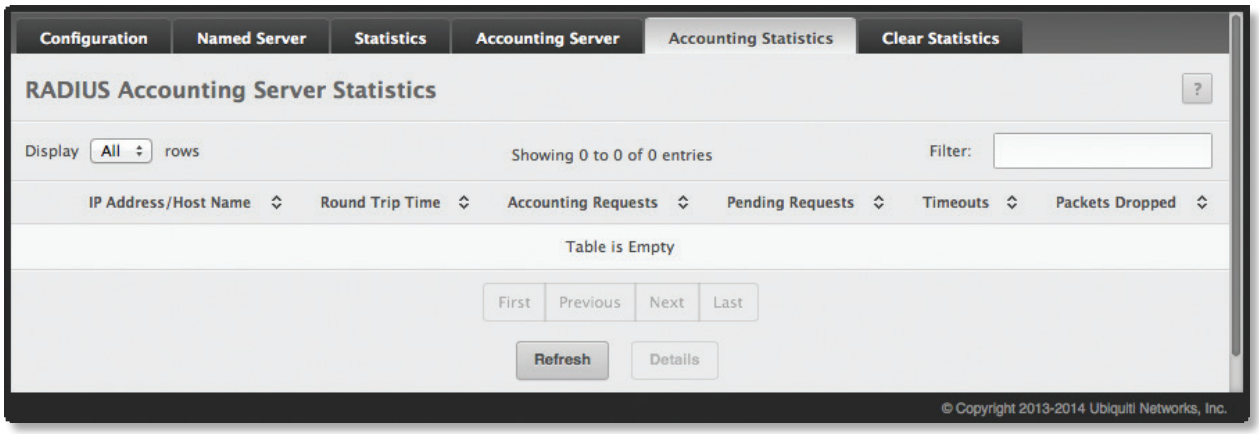
- To add a RADIUS accounting server to the list of servers the RADIUS client can contact, click **Add**, and click **Submit** to apply the changes.
- To change the settings for a configured RADIUS accounting server, select the entry to modify, click **Edit**, configure the settings as needed (you cannot 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
IP Address/Host Name	The IP address or host name of the RADIUS accounting server associated with the rest of the data in the row. When viewing the detailed statistics for a RADIUS accounting server, this field identifies the server.
Round Trip Time	Displays the time interval, in hundredths of a second, between the most recent Accounting-Response and the Accounting-Request that matched it from this RADIUS accounting server.
Accounting Requests	The number of RADIUS Accounting-Request packets sent to this server. This number does not include retransmissions.
Pending Requests	The number of RADIUS Accounting-Request packets destined for the server that have not yet timed out or received a response.
Timeouts	The number of times a response was not received from the server within the configured timeout value.
Packets Dropped	The number of RADIUS packets received from the server on the accounting port and dropped for some other reason.
Accounting Retransmissions	The number of RADIUS Accounting-Request packets retransmitted to the server.
Accounting Responses	The number of RADIUS packets received on the accounting port from the server.
Timeouts	The number of accounting timeouts to this server.
Malformed Access Responses	The number of malformed RADIUS Accounting-Response packets received from the server. Malformed packets include packets with an invalid length. Bad authenticators and unknown types are not included as malformed accounting responses.
Bad Authenticators	The number of RADIUS Accounting-Response packets that contained invalid authenticators received from the accounting server.
Unknown Types	The number of RADIUS packets of unknown type which were received from the server on the accounting port.

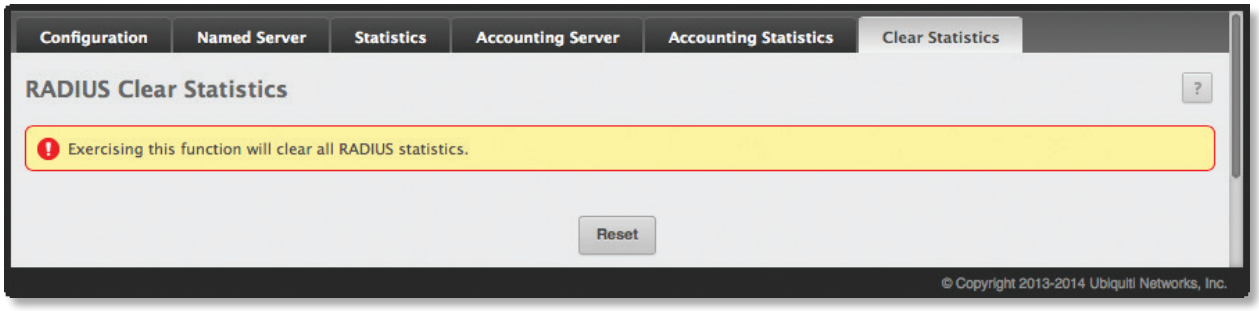
Use the buttons to perform the following tasks:

- 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.



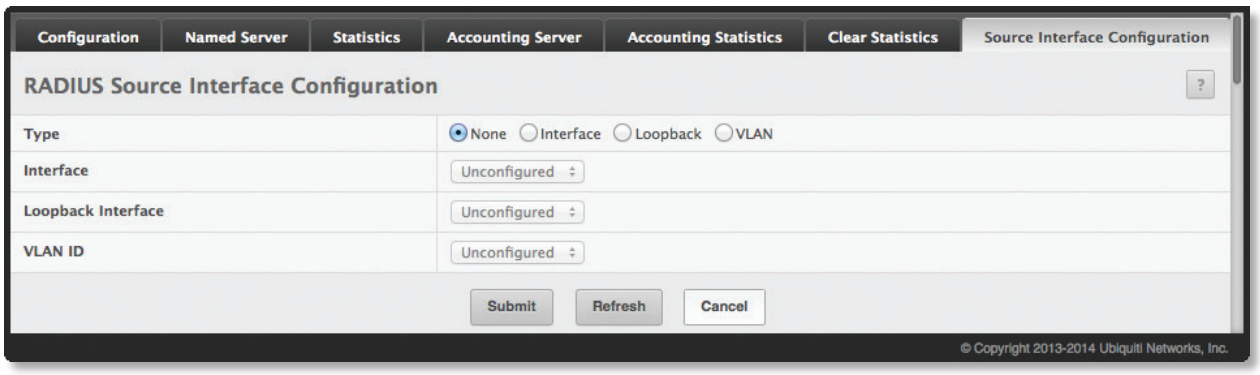
RADIUS Clear Statistics

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 *Interface*, select the physical port to use as the source interface.
VLAN ID	When the selected *Type* is *VLAN*, select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.

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.

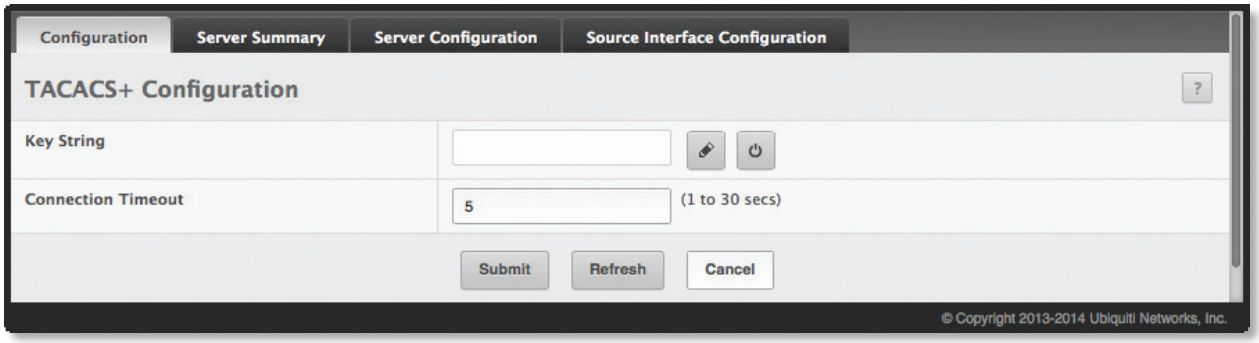
To retain the changes across the switch's next power cycle, click **System** > **Configuration Storage** > **Save**.

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
Key String	Specifies the authentication and encryption key for TACACS+ communications between the device and the TACACS+ server. The key must match the key configured on the TACACS+ server. Click this button to configure the field. Click this button to reset the field to the default value.
Connection Timeout	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:

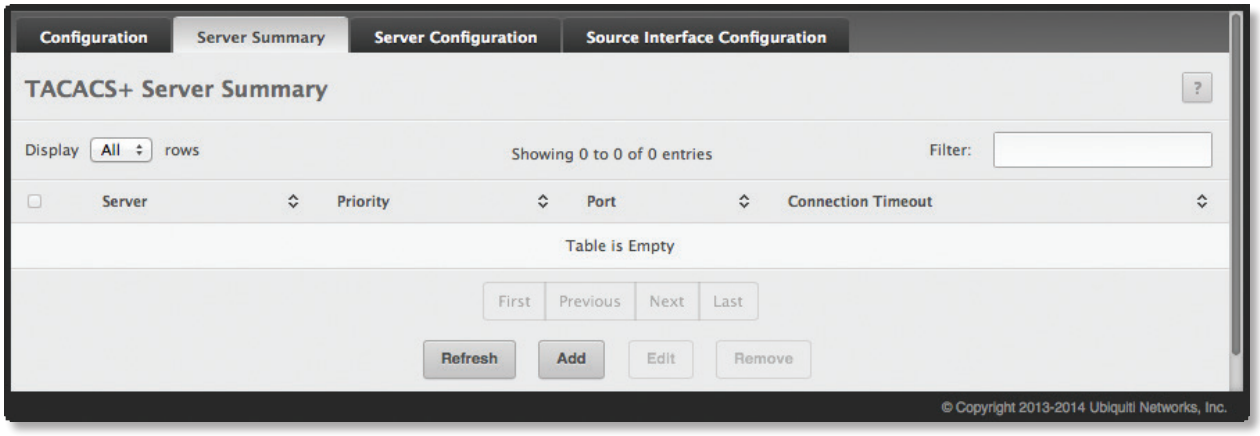
- If you change any settings on this page, click **Submit** to apply the changes.
- Click **Refresh** to refresh the page with the most current data from the switch.

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
Server	Specifies the TACACS+ Server IP address or Hostname.
Priority	Specifies the order in which the TACACS+ servers are used.
Port	Specifies the authentication port.
Connection Timeout	The amount of time that passes before the connection between the device and the TACACS+ server times out.
Add TACACS+ Server dialog box – 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:	
Key String	Specifies the authentication and encryption key for TACACS+ communications between the device and the TACACS+ server. The key must match the encryption used on the TACACS+ server.

Use the buttons to perform the following tasks:

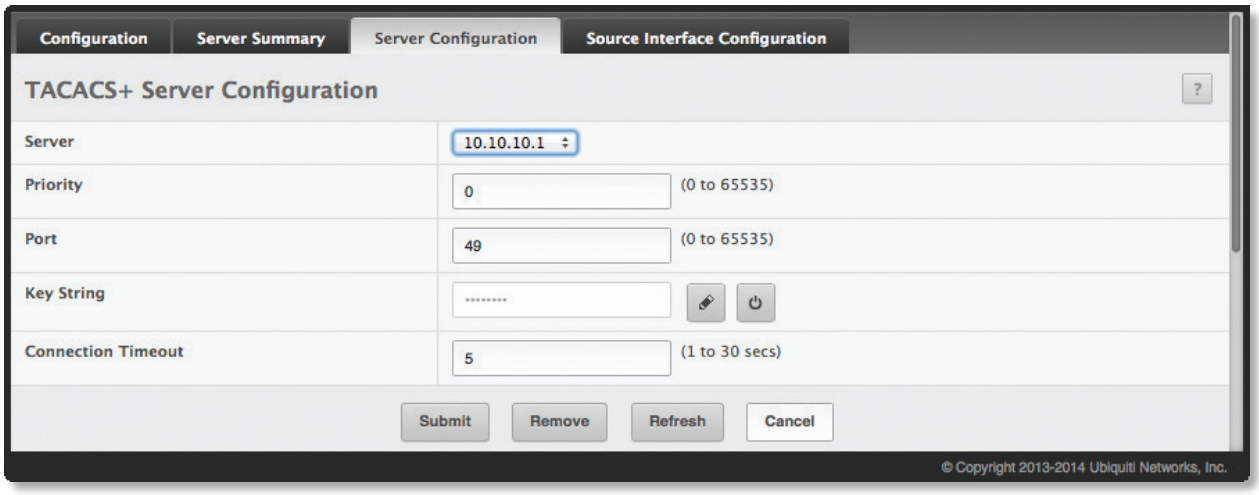
- To add a TACACS+ server to the list of servers the TACACS+ client can contact, click **Add** (this button is disabled if the maximum number of servers has already been added). Configure the fields shown in the table below and click **Submit** to apply the changes.
- To edit a configured TACACS+ server 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
Server	Specifies the TACACS+ Server IP address or Hostname.
Priority	Specifies the order in which the TACACS+ servers are used.
Port	Specifies the authentication port.
Key String	Specifies the authentication and encryption key for TACACS+ communications between the device and the TACACS+ server. The key must match the encryption used on the TACACS+ server.
Connection Timeout	The amount of time that passes before the connection between the device and the TACACS+ server timeout.

Use the buttons to perform the following tasks:

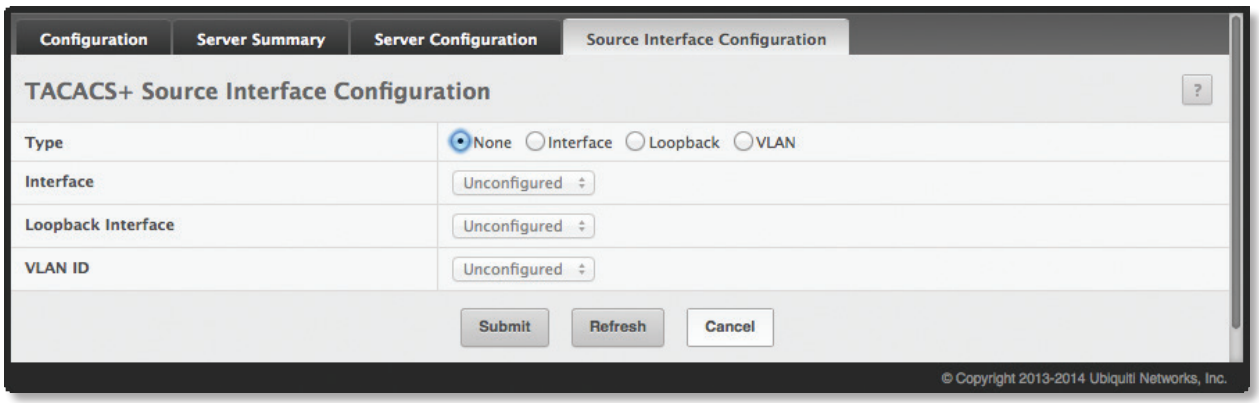
- If you change any fields on this page, click **Submit** to apply the changes.
- To remove a configured TACACS+ server, select it from the table, click **Remove**, and confirm the deletion.
- Click **Refresh** to refresh the page with the most current data from the switch.

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** > T**ACACS+** > **Source Interface Configuration** in the navigation menu.



TACACS+ Source Interface Configuration

TACACS+ 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 *Interface*, select the physical port to use as the source interface.
VLAN ID	When the selected *Type* is *VLAN*, select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.

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.

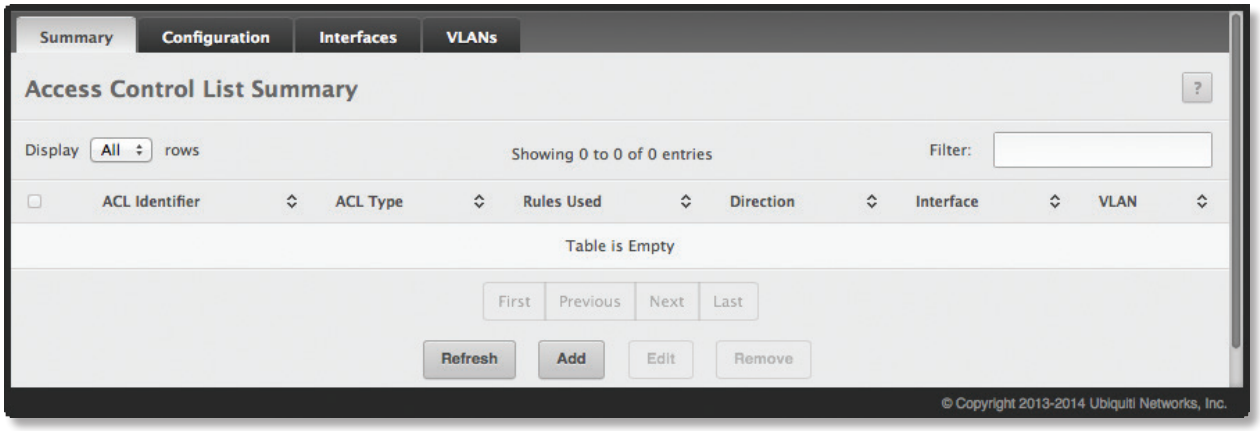
To configure an IP ACL:

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
ACL Identifier	The name or number that identifies the ACL. The permitted identifier depends on the ACL type. Standard and Extended IPv4 ACLs use numbers within a set range, and Named IPv4 and MAC ACLs use alphanumeric characters.
ACL Type	The type of ACL. The ACL type determines the criteria that can be used to match packets. The type also determines which attributes can be applied to matching traffic. IPv4 ACLs classify Layer-3 and Layer-4 IPv4 traffic, IPv6 ACLs classify Layer-3 and Layer-4 IPv6 traffic, and MAC ACLs classify Layer-2 traffic. The ACL types are as follows: • **IPv4 Standard** Match criteria is based on the source address of IPv4 packets. • **IPv4 Extended** Match criteria can be based on the source and destination addresses, source and destination Layer-4 ports, and protocol type of IPv4 packets. • **IPv4 Named** Match criteria is the same as IPv4 Extended ACLs, but the ACL ID can be an alphanumeric name instead of a number. • **IPv6 Named** Match criteria can be based on information including the source and destination IPv6 addresses, source and destination Layer-4 ports, and protocol type within IPv6 packets. • **Extended MAC** Match criteria can be based on the source and destination MAC addresses, 802.1p user priority, VLAN ID, and EtherType value within Ethernet frames.
Rules Used	The number of rules currently configured for the ACL.
Direction	Indicates whether the packet is checked against the rules in an ACL when it is received on an interface (*Inbound*) or after it has been received, routed, and is ready to exit an interface (*Outbound*).
Interface	The interface(s) to which the ACL has been applied.
VLAN	Each VLAN to which the ACL has been applied.

Use the buttons at the bottom of the page to perform the following tasks:

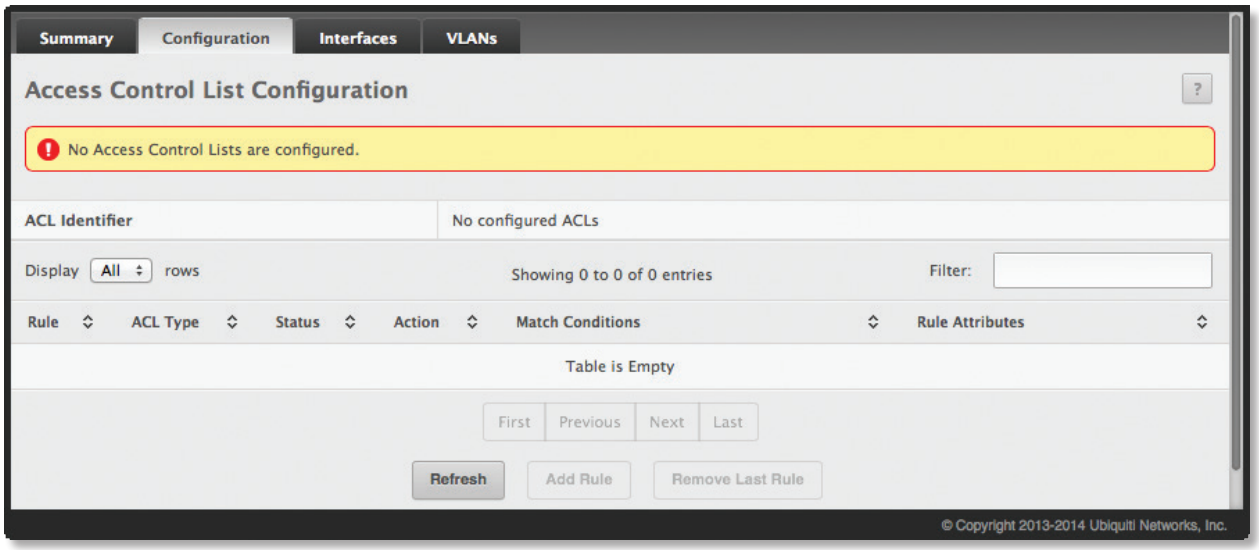
- To add an ACL, click **Add**, configure the ACL type and ID, and click **Submit** to apply the changes.
- To configure rules for an ACL, select the ACL and click **Edit**. Configure the fields on the *Access Control List Configuration* page for the selected ACL (see **"Access Control List Configuration" 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
ACL Identifier	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.
Rule	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.
ACL Type	The type of ACL. The ACL type determines the criteria that can be used to match packets. The type also determines which attributes can be applied to matching traffic.
Status	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.
Action	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 *Add Access Control List Rule* window, the selected action determines which fields can be configured. Not all fields are available for both Permit and Deny actions.
Match Conditions	The criteria used to determine whether a packet or frame matches the ACL rule.
Rule Attributes	Each action, beyond the basic Permit and Deny actions, to perform on the traffic that matches the rule.

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 *ACL Identifier* field. The fields available in the window depend on the *ACL Type*. The following information describes the fields in this window. The Match Criteria tables that apply to IPv4 ACLs, IPv6 ACLs, and MAC ACLs are described separately.	
Match Criteria (IPv4 ACLs) – 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.	
Every	When this option is selected, all packets will match the rule and will be either permitted or denied. This option is exclusive to all other match criteria – if *Every* is selected, no other match criteria can be configured. To configure specific match criteria, this option must be cleared.
*Protocol**	The IANA-assigned protocol number to match within the IP packet. You can also specify one of the following keywords: *EIGRP*, *GRE*, *ICMP*, *IGMP*, *IP*, *IPINIP*, *OSPF*, *PIM*, *TCP*, or *UDP*.
*Fragments**	IP ACL rule to match on fragmented IP packets.
Source IP Address / Wildcard Mask	The source port IP address in the packet and source IP wildcard mask (in the second field) to compare to the IP address in a packet header. Wild card masks determines which bits in the IP address are used and which bits are ignored. A wild card mask of *255.255.255.255* indicates that no bit is important. A wildcard of *0.0.0.0* indicates that all of the bits are important. Wildcard masking for ACLs operates differently from a subnet mask. A wildcard mask is in essence the inverse of a subnet mask. With a subnet mask, the mask has ones (1's) in the bit positions that are used for the network address, and has zeros (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.
*Source L4 Port**	The TCP/UDP source port to match in the packet header. The Source L4 Port and Destination L4 port are configurable only if protocol is either *TCP* or *UDP*. *Equal to*, *Not Equal to*, *Greater than*, and *Less than* options are available. For TCP protocol: *BGP*, *Domain*, *Echo*, *FTP*, *FTP-Data*, *HTTP*, *SMTP*, *Telnet*, *WWW*, *POP2*, or *POP3*. For UDP protocol: *Domain*, *Echo*, *NTP*, *RIP*, *SNMP*, *TFTP*, *Time*, or *WHO*.
Destination IP Address / Wildcard Mask	The destination port IP address in the packet and destination IP wildcard mask (in the second field) to compare to the IP address in a packet header. Wild card masks determines which bits in the IP address are used and which bits are ignored. A wild card mask of *255.255.255.255* indicates that no bit is important. A wildcard of *0.0.0.0* indicates that all of the bits are important. Wildcard masking for ACLs operates differently from a subnet mask. A wildcard mask is in essence the inverse of a subnet mask. With a subnet mask, the mask has ones (1's) in the bit positions that are used for the network address, and has zeros (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.
*Destination L4 Port**	The TCP/UDP destination port to match in the packet header. The Source L4 Port and Destination L4 port are configurable only if protocol is either *TCP* or *UDP*. *Equal to*, *Not Equal to*, *Greater than*, and *Less than* options are available. For TCP protocol: *BGP*, *Domain*, *Echo*, *FTP*, *FTP-Data*, *HTTP*, *SMTP*, *Telnet*, *WWW*, *POP2*, or *POP3*. For UDP protocol: *Domain*, *Echo*, *NTP*, *RIP*, *SNMP*, *TFTP*, *Time*, or *WHO*.
*IGMP Type**	IP ACL rule to match on the specified IGMP message type. Available only if the protocol is IGMP.
*ICMP Type **	IP ACL rule to match on the specified ICMP message type. Available only if the protocol is ICMP.
*ICMP Code**	IP ACL rule to match on the specified ICMP message code. Available only if the protocol is ICMP.
*ICMP Message**	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: *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.*
*TCP Flags**	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

*Service Type**	The service type to match in the IP header. The available options are alternate ways to specify a match condition for the same *Service Type* field in the IP header, but each service type uses a different user notation. After you select the service type, specify the value for the service type in the appropriate field. Only the field associated with the selected service type can be configured. The services types are:
	• **IP DSCP** Matches the packet IP DiffServ Code Point (DSCP) value to the rule. The DSCP value is defined as the high-order six bits of the Service Type octet in the IP header.
	• **IP Precedence** Matches the IP Precedence value to the rule. The IP Precedence field in a packet is defined as the high-order three bits of the Service Type octet in the IP header.
	• **IP TOS Bits** Matches on the Type of Service (TOS) bits in the IP header. The IP TOS field in a packet is defined as all eight bits of the Service Type octet in the IP header. For example, to check for an IP TOS value having bits 7 and 5 set and bit 1 clear, where bit 7 is most significant, use a TOS Bits value of 0xA0 and a TOS Mask of 0xFF.
	  • **TOS Bits** Requires the bits in a packet's TOS field to match the two-digit hexadecimal number entered in this field.
	  • **TOS Mask** The bit positions that are used for comparison against the IP TOS field in a packet.

Time Range Name	The name of the time range that will impose a time 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.
Committed Rate / Burst Size	The allowed transmission rate for packets on the interface (Committed Rate), and the number of bytes allowed in a temporary traffic burst (Burst Rate).
Match Criteria (IPv6 ACLs) – The fields in this section specify the criteria to use to determine whether an IP packet matches the rule. The fields described below apply to IPv6 ACLs.	
Every	When this option is selected, all packets will match the rule and will be either permitted or denied. This option is exclusive to all other match criteria – if *Every* is selected, no other match criteria can be configured. To configure specific match criteria, this option must be cleared.
Protocol	The IANA-assigned protocol number to match within the IP packet. You can also specify one of the following keywords: *ICMP*, *IGMP*, *TCP*, *UDP*, *ICMPv6*, or *IP*.
Fragments	IPv6 ACL rule to match on fragmented IP packets.
Source Prefix / Prefix Length	The IPv6 prefix combined with IPv6 prefix length of the network or host sending the packet.
Source L4 Port	The TCP/UDP source port to match in the packet header. Select one of the following options: *Equal*, *Not Equal*, *Less Than*, *Greater Than*, or *Range*, and specify the port number or keyword. TCP port keywords include *BGP*, *Domain*, *Echo*, *FTP*, *FTP Data*, *HTTP*, *SMTP*, *Telnet*, *WWW*, *POP2*, and *POP3*. UDP port keywords include *Domain*, *Echo*, *NTP*, *RIP*, *SNMP*, *TFTP*, *TIME*, and *WHO*.
Destination Prefix / Prefix Length	The IPv6 prefix combined with the IPv6 prefix length to be compared to a packet's destination IPv6 address as a match criteria for the IPv6 ACL rule. To indicate a destination host, specify an IPv6 prefix length of *128*.
Destination L4 Port	The TCP/UDP destination port to match in the packet header. Select one of the following options: *Equal*, *Not Equal, Less Than, Greater Than*, or *Range*, and specify the port number or keyword. TCP port keywords include *BGP, Domain, Echo, FTP, FTP Data, HTTP, SMTP, Telnet, WWW, POP2*, and *POP3*. UDP port keywords include *Domain, Echo, NTP, RIP, SNMP, TFTP, TIME*, and *WHO*.
ICMP Type	IPv6 ACL rule to match on the specified ICMP message type. This option is available only if the protocol is ICMPv6.
ICMP Code	IPv6 ACL rule to match on the specified ICMP message code. This option is available only if the protocol is ICMPv6.
ICMP Message	IPv6 ACL rule to match on the ICMP message type and code. Specify one of the following supported ICMPv6 messages: *Destination-Unreachable, Echo-Request, Echo-Reply, Header, Hop-Limit, MLD-Query, MLD-Reduction, MLD-Report, ND-NA, ND-NS, Next-Header, No-Admin, No-Route, Packet-Too-Big, Port-Unreachable, Router-Solicitation, Router-Advertisement, Router-Renumbering, Time-Exceeded*, and *Unreachable*. This option is available only if the protocol is ICMPv6.
TCP Flags	IPv6 ACL rule to match on the TCP flags. When a + flag is specified, a match occurs if the flag is set in the TCP header. When a - flag is specified, a match occurs if the flag is not set in the TCP header. When Established is specified, a match occurs if either RST or ACK bits are set in the TCP header. This option is available only if the protocol is TCP.
Flow Label	A 20-bit number that is unique to an IPv6 packet, used by end stations to signify quality-of-service handling in routers.

Access Control List Configuration Fields (Continued)

Field	Description
IP DSCP	The IP DSCP value in the IPv6 packet to match to the rule. The DSCP value is defined as the high-order six bits of the Service Type octet in the IPv6 header.
Routing	IPv6 ACL rule to match on routed packets.
Match Criteria (MAC ACLs) – The fields in this section specify the criteria to use to determine whether an Ethernet frame matches the rule. The fields described below apply to MAC ACLs.	
Every	When this option is selected, all packets will match the rule and will be either permitted or denied. This option is exclusive to all other match criteria – if *Every* is selected, no other match criteria can be configured. To configure specific match criteria, this option must be cleared.
CoS	The 802.1p user priority value to match within the Ethernet frame.
Ethertype	The EtherType value to match in an Ethernet frame. Specify the number associated with the EtherType or specify one of the following keywords: *AppleTalk*, *ARP*, *IBM SNA*, *IPv4*, *IPv6*, *IPX*, *MPLS*, *Unicast*, *NETBIOS*, *NOVELL*, *PPPoE*, or *RARP*.
Source MAC Address / Mask	The MAC address to match to an Ethernet frame's source port MAC address. If desired, enter the MAC mask associated with the source MAC to match. The MAC address mask specifies which bits in the source MAC to compare against an Ethernet frame, 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 *aa:bb:cc:dd:ee:ff*, and the mask is *00:00:ff:ff:ff:ff,* all MAC addresses with *aa:bb:xx:xx:xx:xx* result in a match (where *x* is any hexadecimal number).
Destination MAC Address / Mask	The MAC address to match to an Ethernet frame's destination port MAC address. If desired, enter the MAC Mask associated with the destination MAC to match. The MAC address mask specifies which bits in the destination MAC to compare against an Ethernet frame. Use F's and 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 *aa:bb:cc:dd:ee:ff*, and the mask is *00:00:ff:ff:ff:ff*, all MAC addresses with *aa:bb:xx:xx:xx:xx* result in a match (where *x* is any hexadecimal number).
VLAN	The VLAN ID to match within the Ethernet frame.
Rule Attributes – The fields in this section provide information about the actions to take on a frame or packet that matches the rule criteria. The attributes specify actions other than the basic Permit or Deny actions.	
Assign Queue	The number that identifies the hardware egress queue that will handle all packets matching this rule.
Interface	The interface to use for the action: • **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.
Log	When this option is selected, logging is enabled for this ACL rule (subject to resource availability in the device). If the Access List Trap Flag is also enabled, this will cause periodic traps to be generated indicating the number of times this rule went into effect during the current report interval. A fixed 5 minute report interval is used for the entire system. A trap is not issued if the ACL rule hit count is zero for the current interval.
Time Range Name	The name of the time range that will impose a time limitation on the ACL rule. If a time range with the specified name does not exist, and the ACL containing this ACL rule is associated with an interface, the ACL rule is applied immediately. If a time range with specified name exists, and the ACL containing this ACL rule is associated with an interface, the ACL rule is applied when the specified time-range becomes active. The ACL rule is removed when the specified time-range with becomes inactive.
Committed Rate / Burst Size	The allowed transmission rate for frames on the interface (Committed Rate), and the number of bytes allowed in a temporary traffic burst (Burst Rate).

Use the buttons to perform the following tasks:

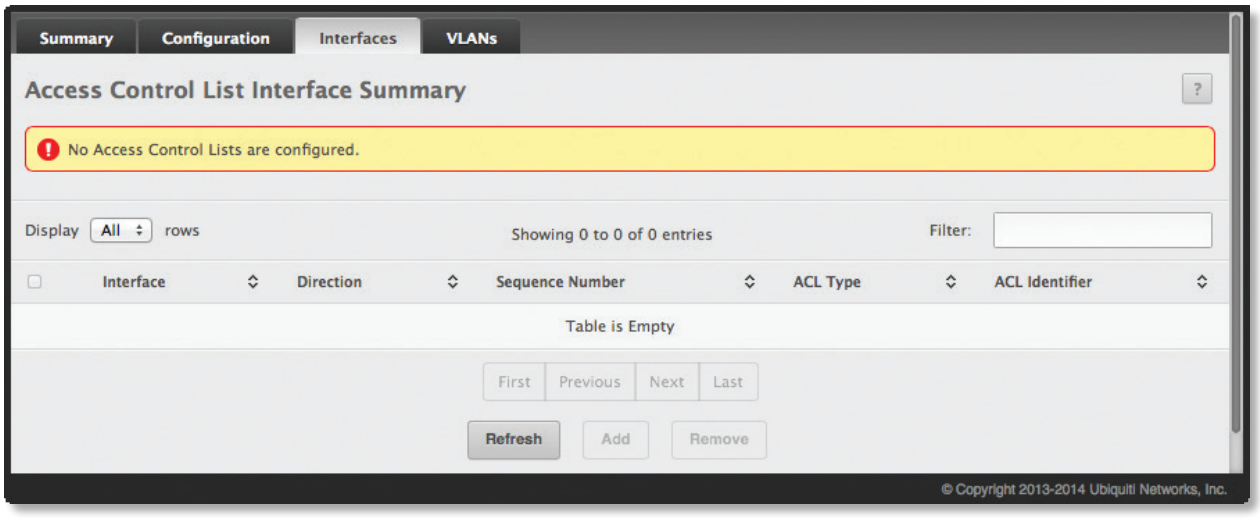
- To add an ACL rule entry, select the ID of the ACL that will include the rule from the *ACL Identifier* drop-down menu. Then, click **Add Rule** and configure the rule criteria and attributes (new rules cannot be created if the maximum number of rules has been reached). Finally, click **Submit** to 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
Interface	The interface that has an associated ACL.
Direction	Indicates whether the packet is checked against the rules in an ACL when it is received on an interface (*Inbound*) or after it has been received, routed, and is ready to exit an interface (*Outbound*).
Sequence Number	The order the ACL is applied to traffic on the interface relative to other ACLs associated with the interface in the same direction. When multiple ACLs are applied to the same interface in the same direction, the ACL with the lowest sequence number is applied first, and the other ACLs are applied in ascending numerical order.
ACL Type	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.
ACL Identifier	The name or number that identifies the ACL. When applying an ACL to an interface, the *ACL Identifier* menu includes only the ACLs within the selected *ACL Type*.

Use the buttons to perform the following tasks:

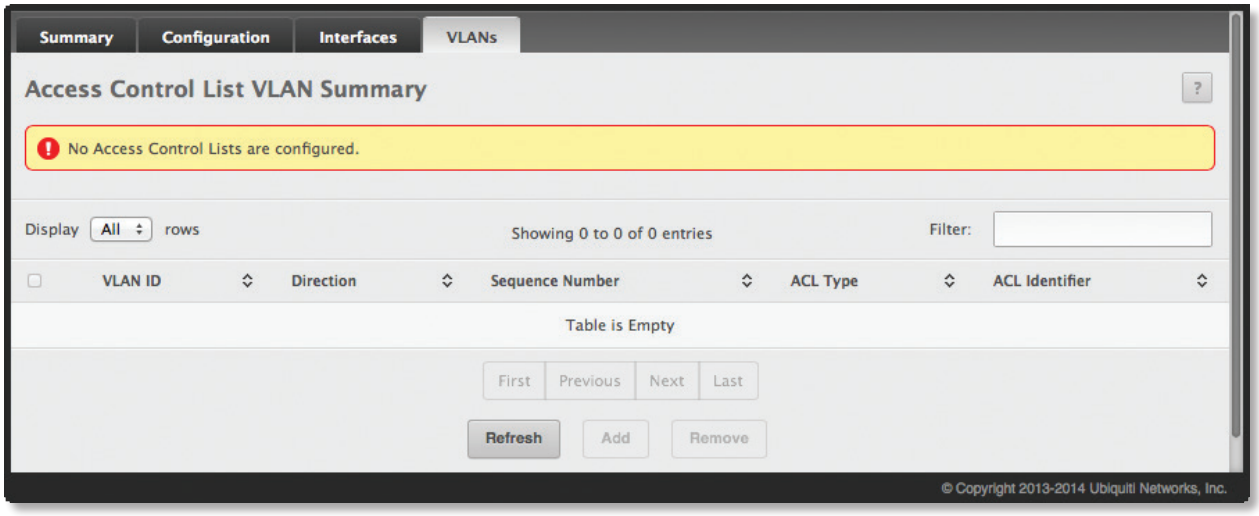
- To apply an ACL to an interface, click **Add** and configure the settings in the available fields.
- To remove the association between an interface and an ACL, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.
- Click **Refresh** to refresh the 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
VLAN ID	The ID of the VLAN associated with the rest of the data in the row. When associating a VLAN with an ACL, use this field to select the desired VLAN.
Direction	Indicates whether the packet is checked against the rules in an ACL when it is received on a VLAN (*Inbound*) or after it has been received, routed, and is ready to exit a VLAN (*Outbound*).
Sequence Number	The order the ACL is applied to traffic on the VLAN relative to other ACLs associated with the VLAN in the same direction. When multiple ACLs are applied to the same VLAN in the same direction, the ACL with the lowest sequence number is applied first, and the other ACLs are applied in ascending numerical order.
ACL Type	The 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.
ACL Identifier	The name or number that identifies the ACL. The permitted identifier depends on the ACL type. Standard and Extended IPv4 ACLs use numbers within a set range, and Named IPv4, IPv6, and MAC ACLs use alphanumeric characters.

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:

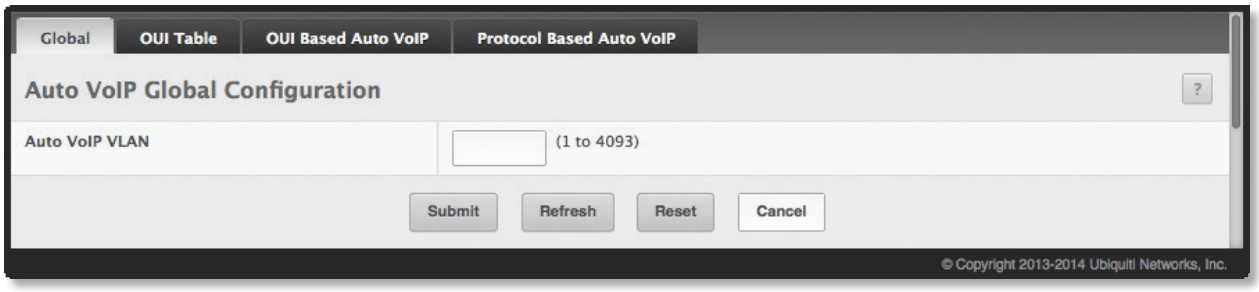
- Session Initiation Protocol (SIP)
- H.323
- Skinny Client Control Protocol (SCCP)

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
Auto VoIP VLAN	The VLAN used to segregate VoIP traffic from other non-voice traffic.

Use the buttons to perform the following tasks:

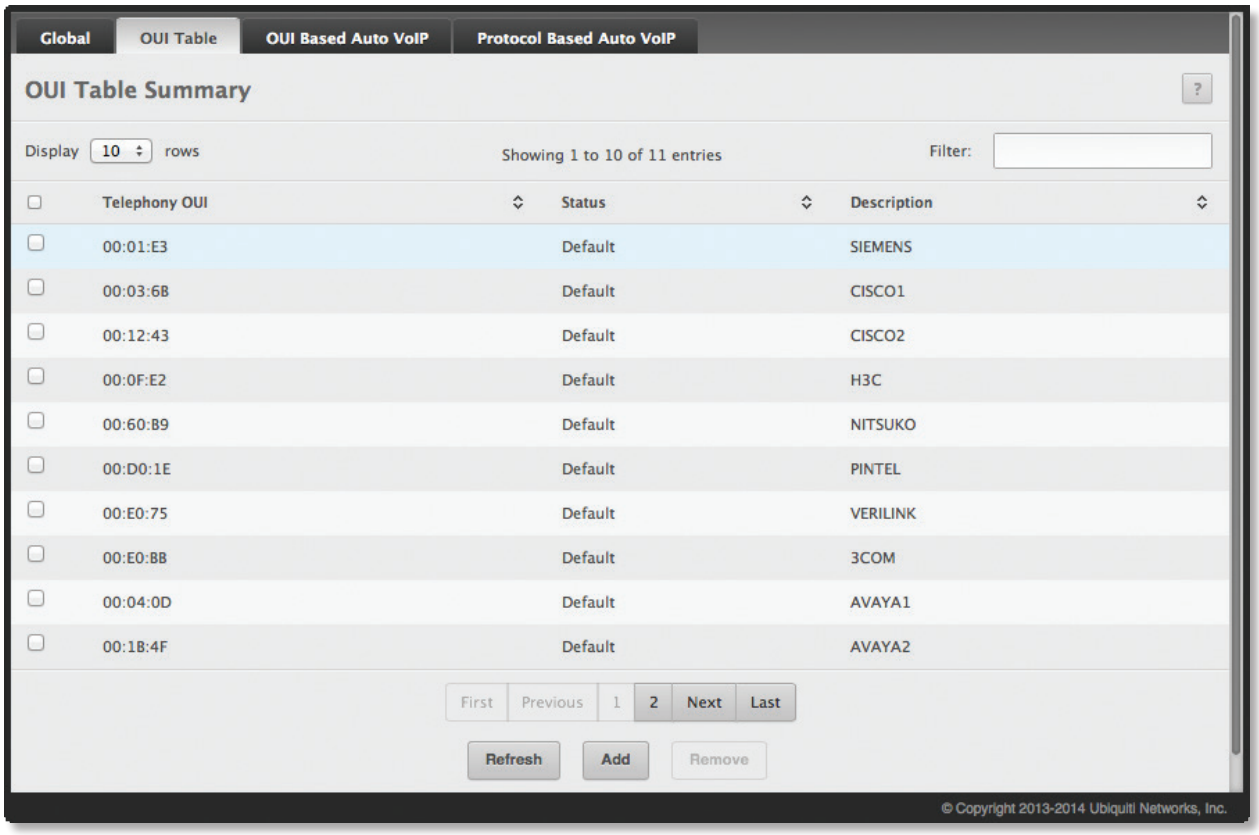
- If you change the *Auto VoIP VLAN* field, click **Submit** to apply the change.
- To reset the VLAN to the default Auto VoIP VLAN, click **Reset**, and confirm the action.
- Click **Refresh** to refresh the page with the most current data from the switch.

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
Telephony OUI	The unique OUI that identifies the device manufacturer or vendor. The OUI is specified in three octet values (each octet is represented as two hexadecimal digits) separated by colons.
Status	Identifies whether the OUI is preconfigured on the system (*Default*) or added by a user (*Configured*).
Description	Identifies the manufacturer or vendor associated with the OUI.

Use the buttons to perform the following tasks:

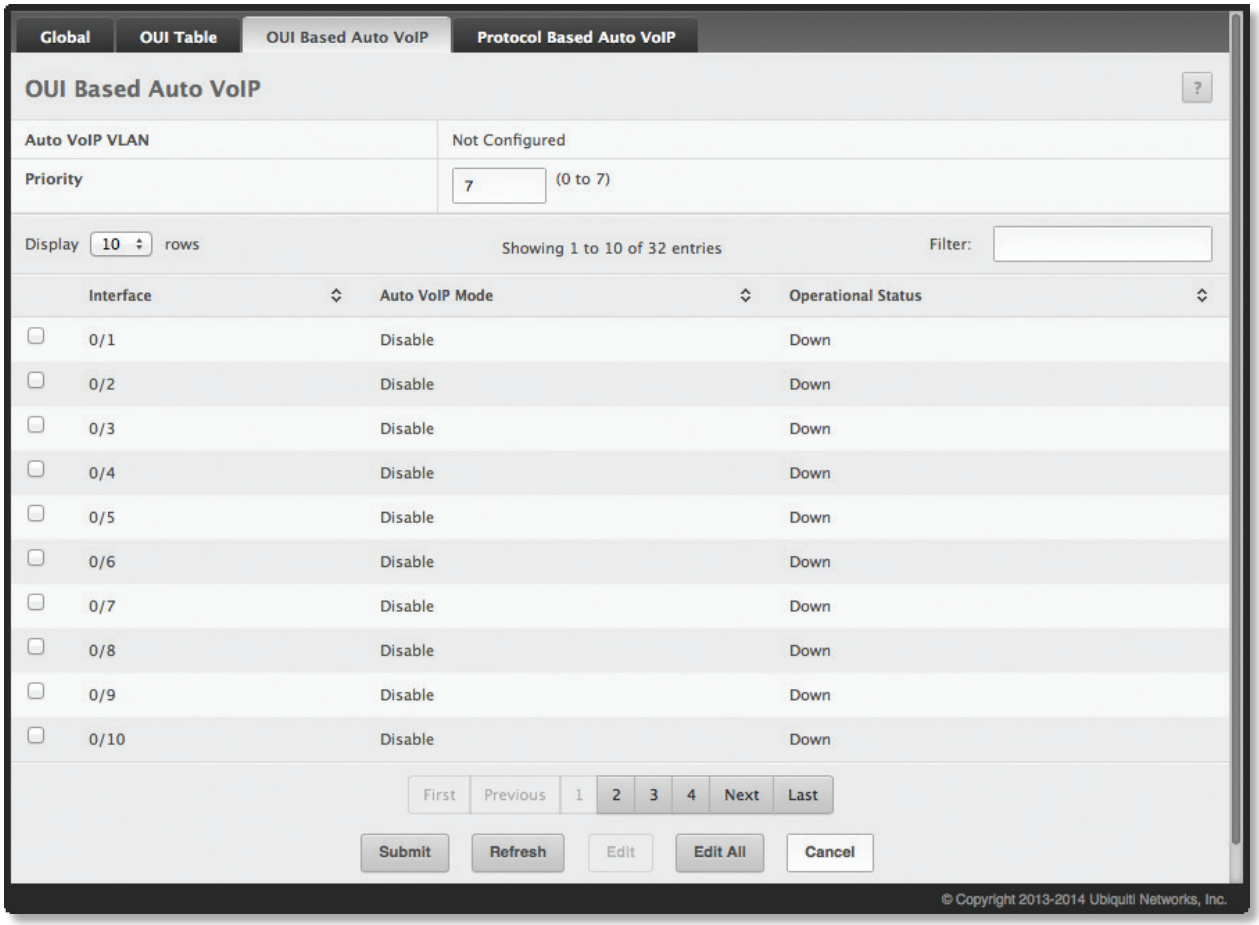
- To add an OUI, click **Add**, specify an OUI and its description in the available fields, and click **Submit** to apply the changes.
- To remove one or more configured OUIs, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.
- Click **Refresh** 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.



OUI Based Auto VoIP

OUI Based Auto VoIP Fields

Field	Description
Auto VoIP VLAN	The VLAN used to segregate VoIP traffic from other non-voice traffic. All VoIP traffic that matches a value in the known OUI list gets assigned to this VoIP VLAN.
Priority	The 802.1p priority used for traffic that matches a value in the known OUI list. If the Auto VoIP mode is enabled and the interface detects an OUI match, the device assigns the traffic in that session to the traffic class mapped to this priority value. Traffic classes with a higher value are generally used for time-sensitive traffic.
Interface	The interface associated with the rest of the data in the row. When editing Auto VoIP settings on one or more interfaces, this field identifies the interface(s) being configured.
Auto VoIP Mode	The administrative mode of OUI-based Auto VoIP on the interface.
Operational Status	The interface's operational status (*Up* or *Down*). To be *Up*, 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.

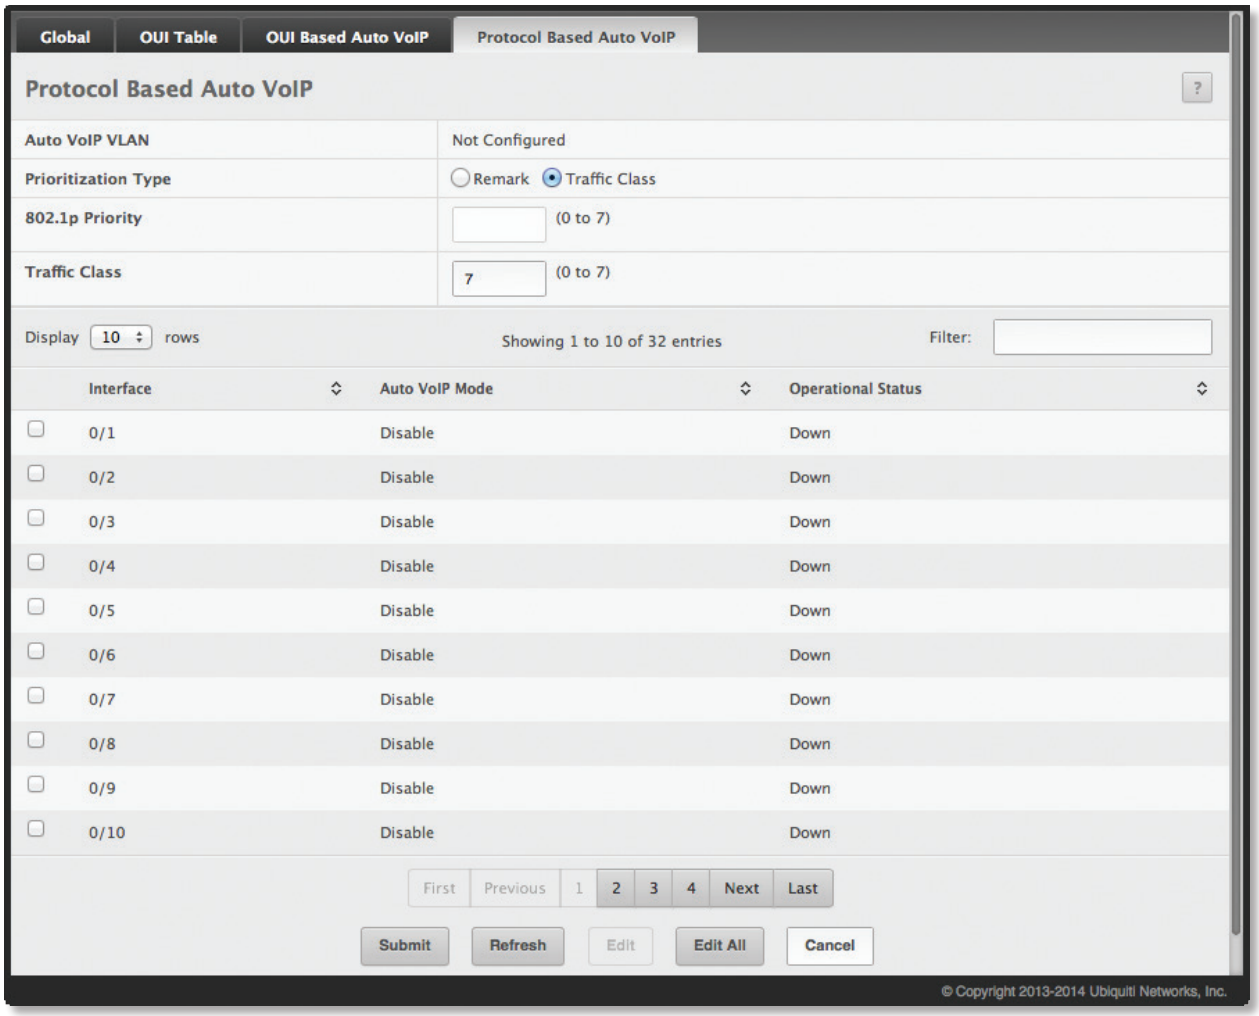
- To configure settings on all interfaces, click **Edit All**. In the *Edit OUI Based Port Configuration* window, change the settings as needed, and click **Submit** to apply the changes.
- Click **Refresh** to refresh the page with the most current data from the switch.

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.



Protocol Based Auto VoIP

Protocol Based Auto VoIP Fields

Field	Description
Auto VoIP VLAN	The VLAN used to segregate VoIP traffic from other non-voice traffic. All VoIP traffic in a session identified by the call-control protocol gets assigned to this VoIP VLAN.
Prioritization Type	The method used to prioritize VoIP traffic when a call-control protocol is detected, which is one of the following: • **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 *Prioritization Type* is 802.1p Priority. If the *Auto VoIP Mode* is enabled and the interface detects a call-control protocol, the device marks traffic in that session with the specified 802.1p priority value to ensure voice traffic always gets the highest priority throughout the network path. Egress tagging must be administratively enabled on the appropriate uplink port to carry the remarked priority at the egress port.
Traffic Class	The traffic class used for protocol-based VoIP traffic. This field can be configured if the *Prioritization Type* is **Traffic Class**. If the *Auto VoIP Mode* is enabled and the interface detects a call-control protocol, the device assigns the traffic in that session to the configured 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

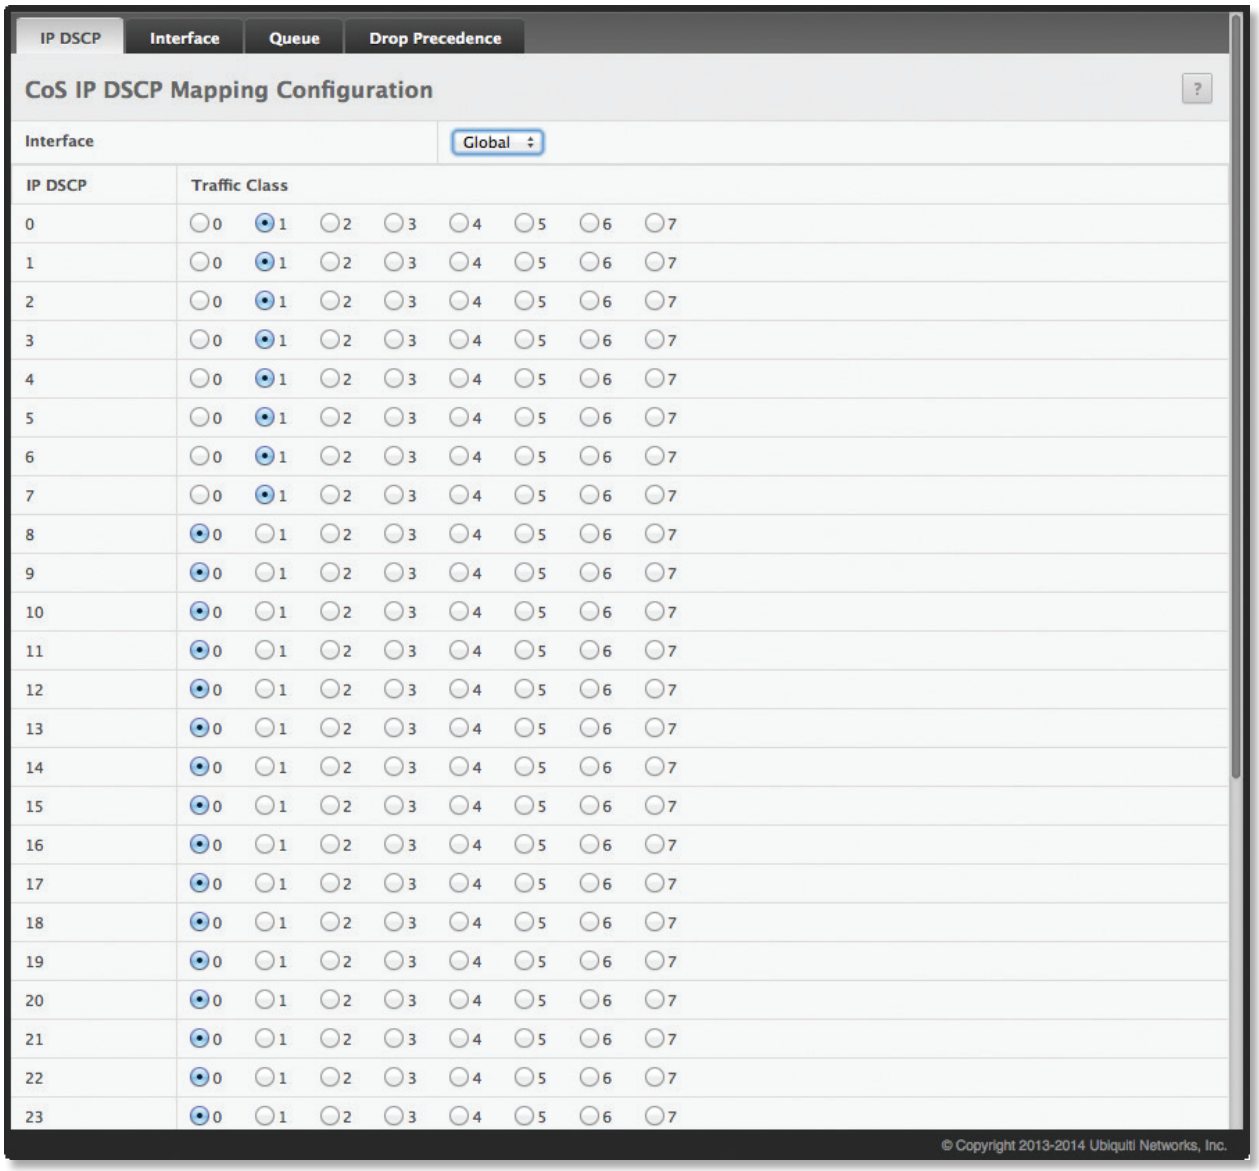
The Class of Service (CoS) queueing feature lets you directly configure certain aspects of switch queueing. This provides the desired QoS behavior for different types of network traffic when the complexities of DiffServ are not required. The priority of a packet arriving at an interface can be used to 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.

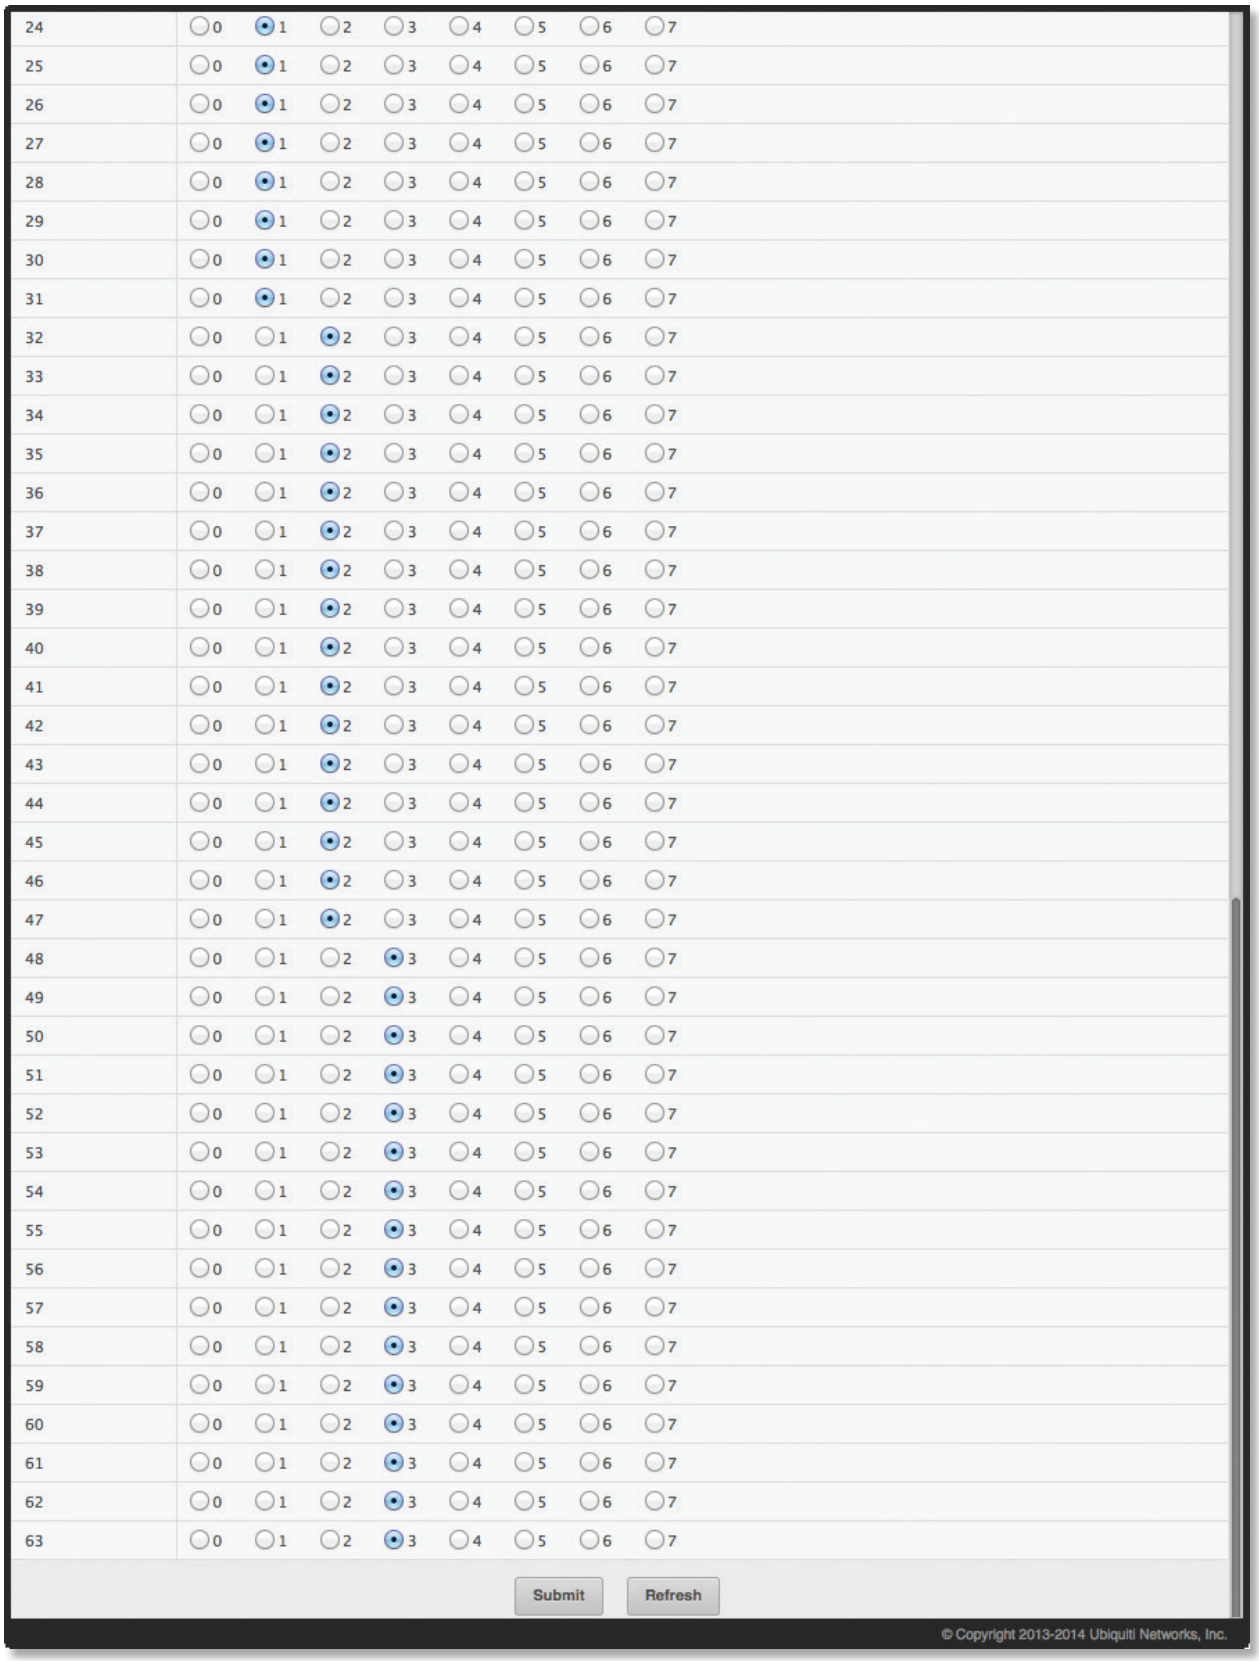
CoS IP DSCP Mapping Configuration

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.



CoS IP DSCP Mapping Configuration – 1 of 2



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:

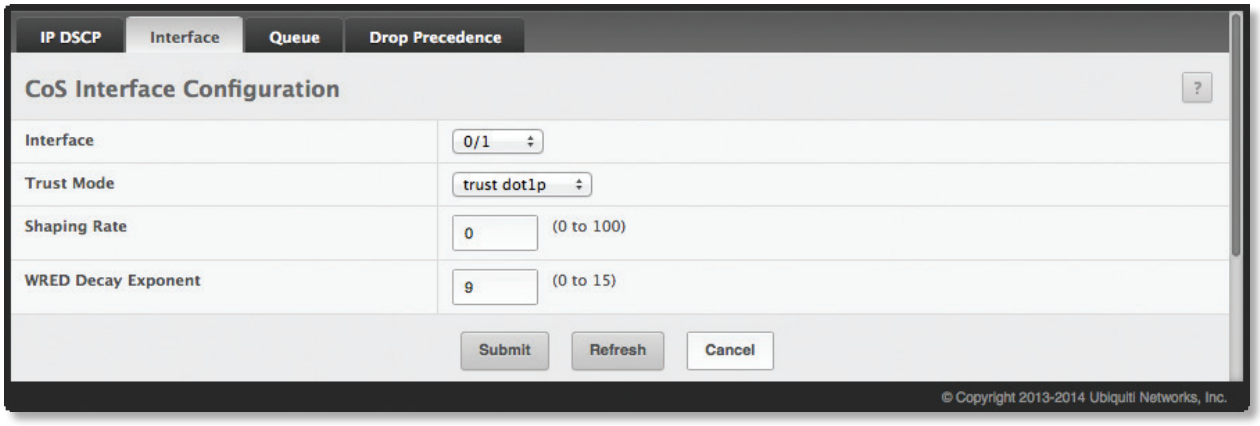
- If you change any fields, click **Submit** to apply the changes.
- Click **Refresh** to refresh the page with the most current data from the switch.

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:

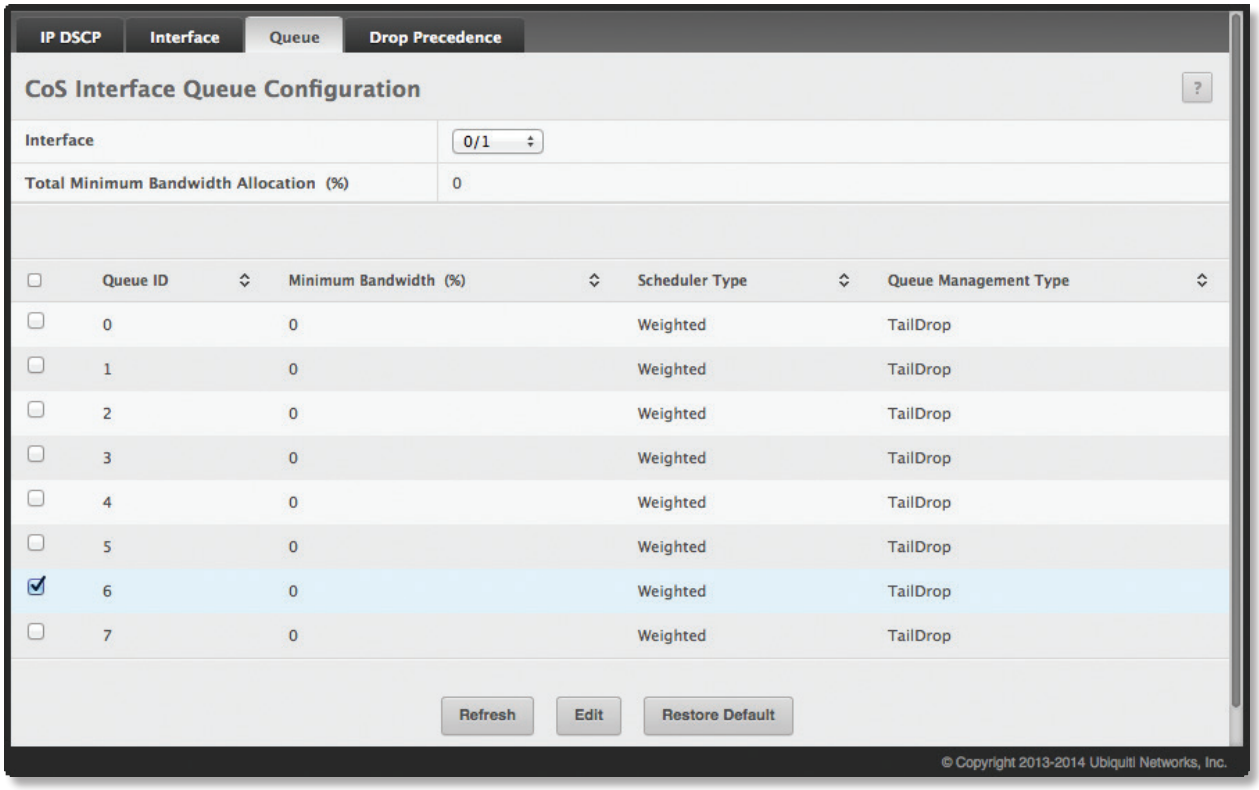
- If you change any fields on the page, click **Submit** to apply the changes.
- Click **Refresh** to refresh the page with the most current data from the switch.

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.



CoS Interface Queue Configuration

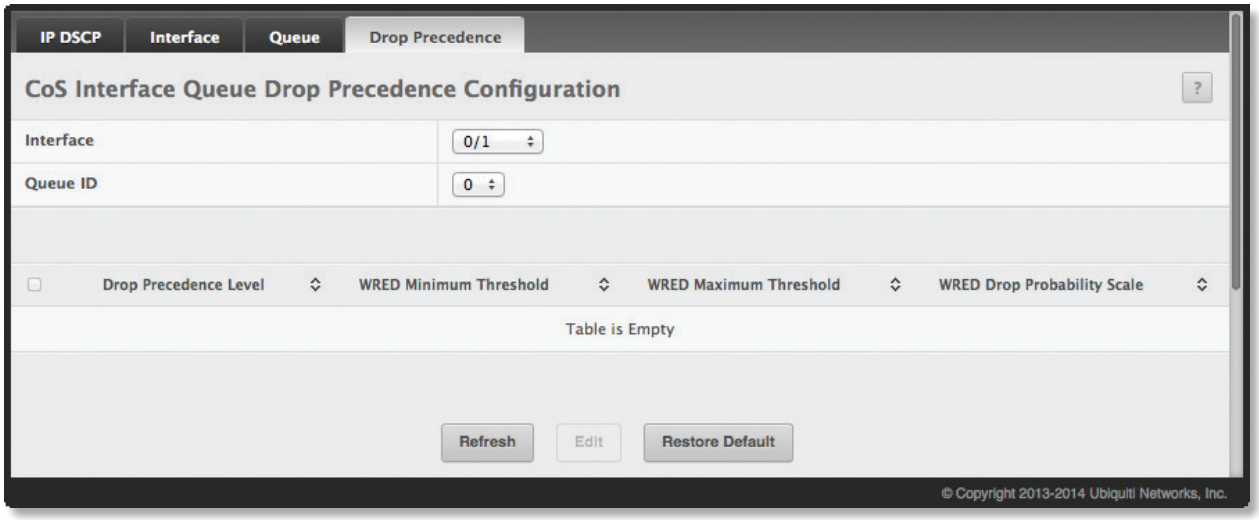
CoS Interface Queue Configuration Fields

Field	Description
Interface	Specifies the interface (physical, LAG, or Global) to configure.
Total Minimum Bandwidth Allocation (%)	Shows the total minimum bandwidth allocated to the selected interface for all the queues.
Queue ID	The CoS queue. The higher the queue value, the higher its priority is for sending traffic.
Minimum Bandwidth	The minimum guaranteed bandwidth allocated to the selected queue on the interface. Setting this value higher than its corresponding *Maximum Bandwidth* automatically increases the maximum to the same value. A value of *0* (zero) means no guaranteed minimum. The sum of individual *Minimum Bandwidth* values for all queues in the selected interface cannot exceed defined a maximum of 100.
Scheduler Type	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.
Queue Management Type	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
Interface	The interface on which to configure the queue drop precedence settings. To configure the same settings on all interfaces, select the *Global* menu option.
Queue ID	The CoS queue on which to configure the drop precedence settings. The higher the queue value, the higher its priority is for sending traffic.
Drop Precedence Level	The four drop precedence levels.
WRED Minimum Threshold	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.
WRED Maximum Threshold	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.
WRED Drop Probability Scale	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

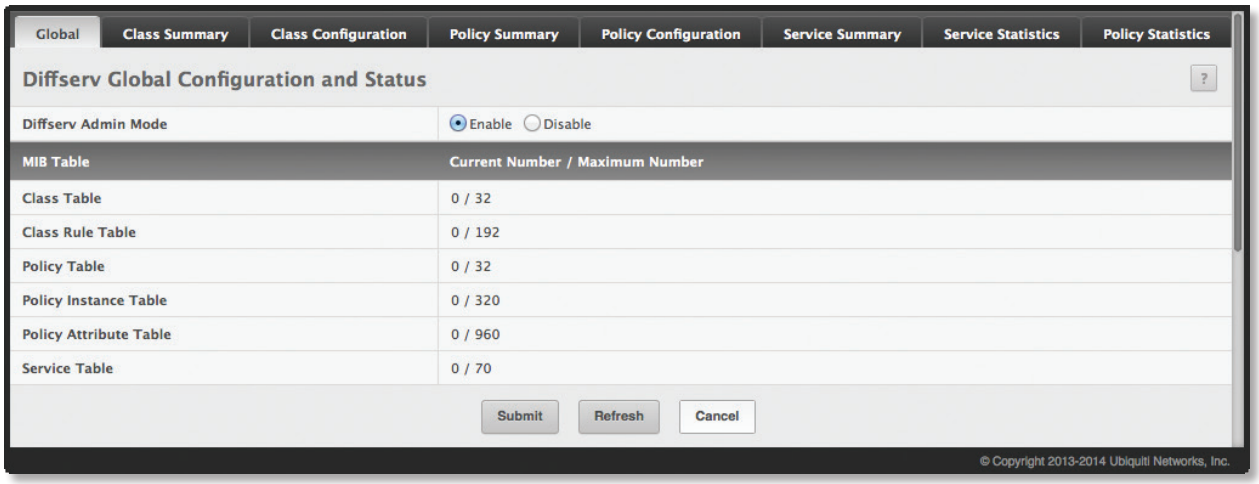
Use this page to configure the administrative mode of Differentiated Services (DiffServ) support on the device and to view the current and maximum number of entries in each of the main DiffServ private MIB tables. DiffServ allows traffic to be classified into streams and given certain QoS treatment in 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
Diffserv Admin Mode	The administrative mode of DiffServ on the device. While disabled, the DiffServ configuration is retained and can be changed, but it is not active. While enabled, Differentiated Services are active.
MIB Table – The information in this table displays the number of entries (rows) that are currently in each of the main DiffServ private MIB tables and the maximum number of rows that can exist in each table.	
Class Table	The current and maximum number of classifier entries in the table. DiffServ classifiers differentiate among traffic types.
Class Rule Table	The current and maximum number of class rule entries in the table. Class rules specify the match criteria that belong to a class definition.
Policy Table	The current and maximum number of policy entries in the table. The policy determines the traffic conditioning or service provisioning actions applied to a traffic class.
Policy Instance Table	The current and maximum number of policy-class instance entries in the table. A policy-class instance is a policy that is associated with an existing DiffServ class.
Policy Attribute Table	The current and maximum number of policy attribute entries in the table. A policy attribute entry attaches various policy attributes to a policy-class instance.
Service Table	The current and maximum number of service entries in the table. A service entry associates a DiffServ policy with an interface and inbound or outbound direction.

Use the buttons to perform the following tasks:

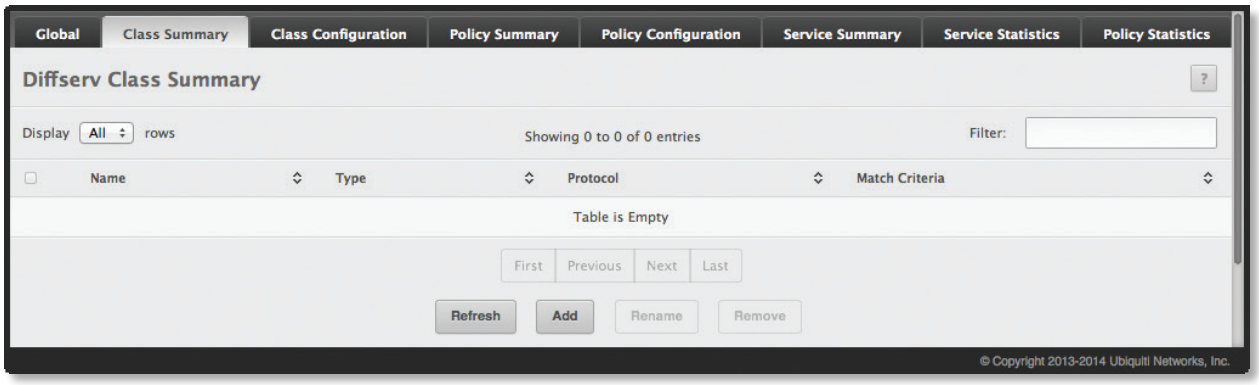
- If you change any fields on this page, click **Submit** to apply the changes.
- Click **Refresh** to refresh the page with the most current data from the switch.

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
Name	The name of the DiffServ class. When adding a new class or renaming an existing class, the name of the class is specified in the *Class* field of the dialog window.
Type	The class type, which is one of the following: • **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.
Protocol	The Layer-3 protocol to use for filtering class types, which is either *IPv4* or *IPv6*.
Match Criteria	The criteria used to match packets.

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

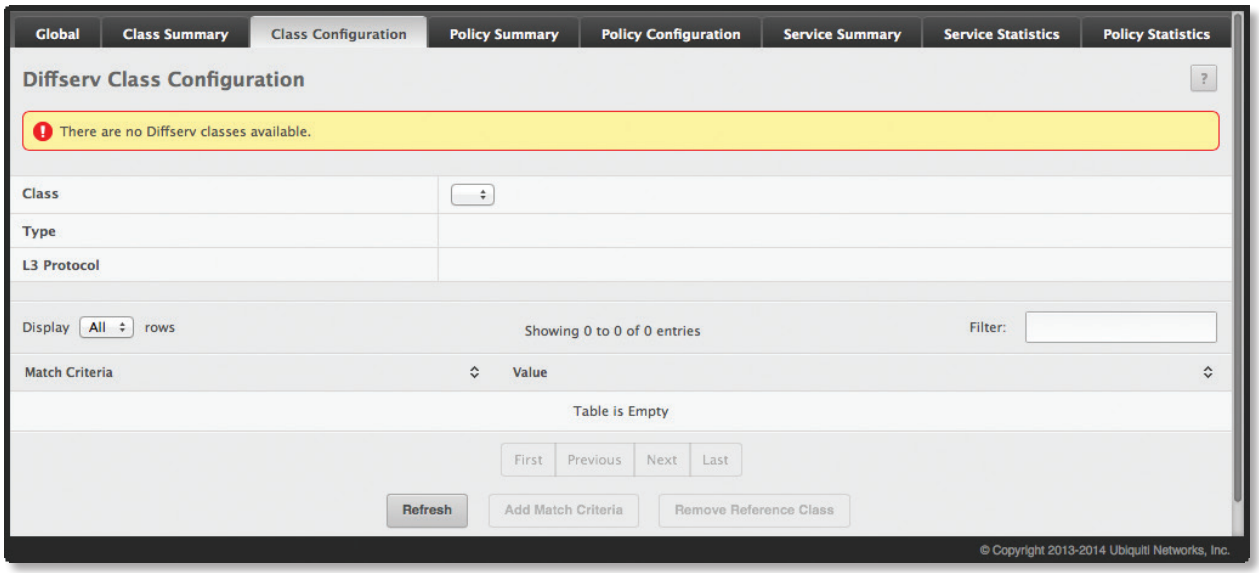
Use this page to define the criteria to associate with a DiffServ class. As packets are received or transmitted, these DiffServ classes are used to classify and prioritize packets. Each class can contain multiple match criteria.

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
Class	The name of the class. To configure match criteria for a class, select its name from the menu.
Type	The class type, which is one of the following: • **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.
Protocol	The Layer-3 protocol to use for filtering class types, which is either IPv4 or IPv6.
Match Criteria	The type of match criteria defined for the selected class. If the *Type* is *ACL*, no information about the match criteria is available on this page.
Value	The configured value of the match criteria that corresponds to the match type.
Add Match Criteria 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.	
Any	Select this option to specify that all packets are considered to match the specified class. There is no need to configure additional match criteria if *Any* is selected because a match will occur on all packets.

Diffserv Class Configuration Fields (Continued)

Field	Description
Reference Class	Select this option to reference another class for criteria. The match criteria defined in the referenced class is as match criteria in addition to the match criteria you define for the selected class. After selecting this option, the classes that can be referenced are displayed. Select the class to reference. A class can reference at most one other class of the same type.
Class of Service	Select this option to require the Class of Service (CoS) value in an Ethernet frame header to match the specified CoS value.
Secondary Class of Service	Select this option to require the secondary CoS value in an Ethernet frame header to match the specified secondary CoS value.
Ethertype	Select this option to require the *EtherType* value in the Ethernet frame header to match the specified *EtherType* value. After you select this option, specify the *EtherType* value in one of these two fields: • *Ethertype Keyword* – The menu includes several common protocols that are mapped to their EtherType values. • *Ethertype Value* – This field accepts custom EtherType values.
VLAN	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.
Secondary VLAN	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.
Source MAC Address	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 *aa:bb:cc:dd:ee:ff*, and the mask is *ff:ff:00:00:00:00*, all MAC addresses with *aa:bb:xx:xx:xx:xx* result in a match (where *x* is any hexadecimal number). Note that this is not a wildcard mask, which ACLs use.
Destination MAC Address	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 aa:bb:cc:dd:ee:ff, and the mask is *ff:ff:00:00:00:00*, all MAC addresses with *aa:bb:xx:xx:xx:xx* result in a match (where *x* is any hexadecimal number). Note that this is not a wildcard mask, which ACLs use.
Source IPv6 Address	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.
Destination IPv6 Address	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
Source L4 Port	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.
Destination L4 Port	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.
IP DSCP	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.
IP Precedence	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.
IP TOS	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.
Protocol	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.
Flow Label	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:

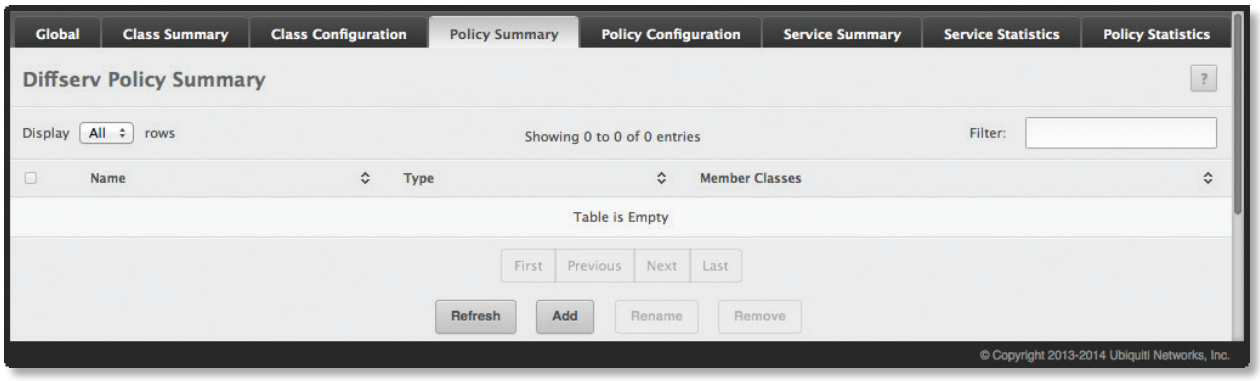
- To define the match criteria for the selected class, click **Add Match Criteria**, In the *Add Match Criteria* window, configure the fields shown in the table below, and click **Submit** to apply the changes. Once you add a match criteria entry to a class, you cannot edit or remove the entry. However, 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
Name	The name of the DiffServ policy. When adding a new policy or renaming an existing policy, the name of the policy is specified in the *Policy* field of the *Add Policy* dialog box.
Type	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.
Member Classes	The DiffServ class or classes that have been added to the policy.

Use the buttons to perform the following tasks:

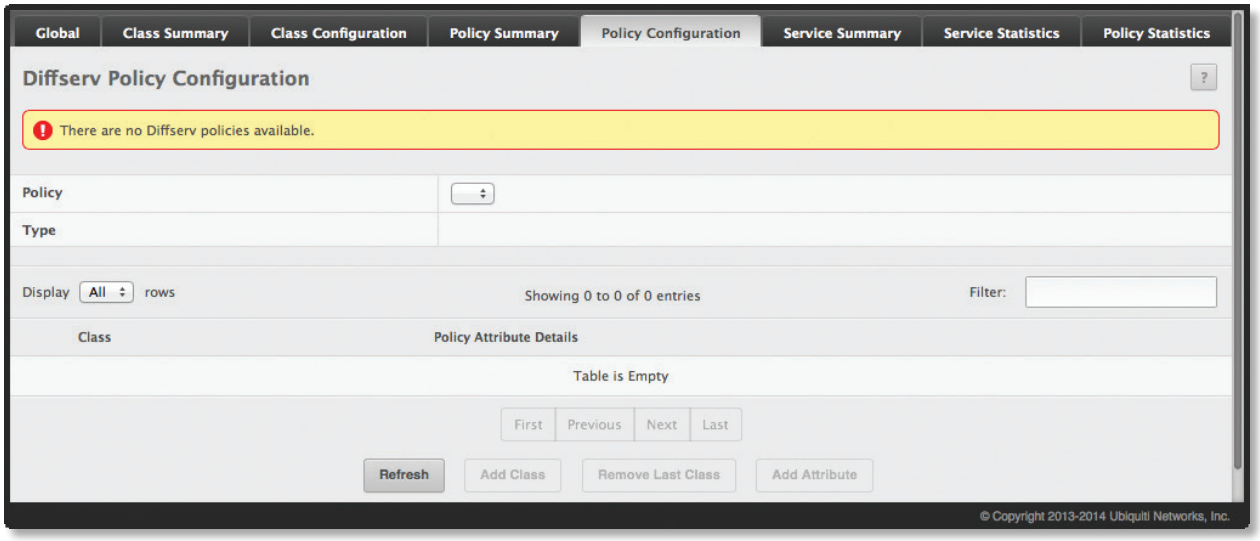
- To add a DiffServ policy, click **Add**, configure the fields, and click **Submit** to apply the changes.
- To change the name of an existing policy, select the entry to modify and click **Rename**, enter the new name, and click **Submit** to apply the change.
- To remove one or more configured policies, 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
Policy	The name of the policy. To add a class to the policy, remove a class from the policy, or configure the policy attributes, you must first select its name from the menu.
Type	The traffic flow direction to which the policy is applied.
Class	The DiffServ class or classes associated with the policy. The policy is applied to a packet when a class match within that policy-class is found.
Policy Attribute Details	The policy attribute types and their associated values that are configured for the policy.
Add Policy Attribute window – 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.	
Assign Queue	Select this option to assign matching packets to a traffic queue. Use the Queue ID Value field to select the queue to which the packets of this policy-class are assigned.
Drop	Select this option to drop packets that match the policy-class.
Mark CoS	Select this option to mark all packets in a traffic stream with the specified Class of Service (CoS) queue value. Use the *Class of Service* field to select the CoS value to mark in the priority field of the 802.1p header (the only tag in a single tagged packet or the first or outer 802.1Q tag of a double VLAN tagged packet). If the packet does not already contain this header, one is inserted.
Mark CoS as Secondary CoS	Select this option to mark all packets in a traffic stream with the specified secondary CoS queue number. Use the *Class of Service* field to select the CoS value to mark in the priority field of the 802.1p header in the secondary (inner) 802.1Q tag of a double VLAN tagged packet. If the packet does not already contain this header, one is inserted.
Mark IP DSCP	Select this option to mark all packets in the associated traffic stream with the specified IP DSCP value. 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.
Mark IP Precedence	Select this option to mark all packets in the associated traffic stream with the specified IP Precedence value. Then, select the *IP Precedence Value* to mark in packets that match the policy-class.
Mirror Interface	Select this option to copy the traffic stream to a specified egress port (physical or LAG) without bypassing normal packet forwarding. This 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
Police Simple	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: *CoS*, *IP DSCP*, *IP Precedence*, or *Secondary COS*. • **Committed Rate (Kbps)** The maximum allowed arrival rate of incoming packets for this class. • **Committed Burst Size (Kbytes)** The amount of conforming traffic allowed in a burst. • **Conform Action** The action taken on packets considered to be conforming (below the police rate). • **Violate Action** The action taken on packets considered to be non-conforming (above the police rate).
Police Single Rate	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: *CoS*, *IP DSCP*, *IP Precedence*, or *Secondary COS*. This field is available only if one or more classes that meets the color-awareness criteria exist. • **Color Exceed Class** For color-aware policing, packets are metered against the PIR only. • **Committed Rate (Kbps)** The maximum allowed arrival rate of incoming packets for this class. • **Committed Burst Size (Kbytes)** The amount of conforming traffic allowed in a burst. • **Excess Burst Size (Kbytes)** The amount of conforming traffic allowed to accumulate beyond the Committed Burst Size (Kbytes) value during longer-than-normal idle times. This value allows for occasional bursting. • **Conform Action** The action taken on packets 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).
Police Two Rate	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: *CoS*, *IP DSCP*, *IP Precedence*, or *Secondary COS*. This field is available only if one or more classes that meets the color-awareness criteria exist. • **Color Exceed Class** For color-aware policing, packets are metered against the PIR. • **Committed Rate (Kbps)** The maximum allowed arrival rate of incoming packets for this class. • **Committed Burst Size (Kbytes)** The amount of conforming traffic allowed in a burst. • **Peak Rate (Kbps)** The maximum 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).
Redirect Interface	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:

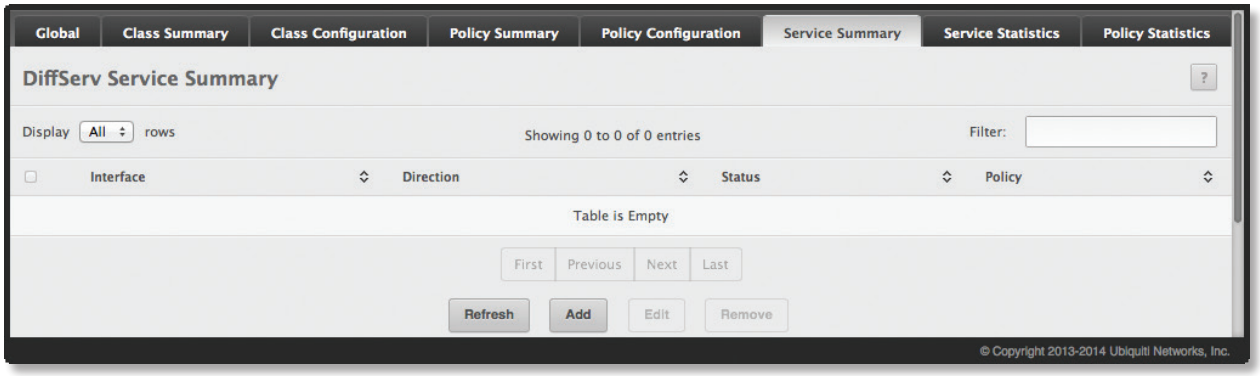
- To add a class to the policy, click **Add Class**.
- To add attributes to a policy or to change the policy attributes, select the policy with the attributes to configure and click **Add Attribute**.
- To remove the most recently associated class from the selected policy, click **Remove Last Class**.
- 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
Interface	The interface associated with the rest of the data in the row. Only interfaces that have an associated policy are listed in the table.
Direction	The traffic flow direction to which the policy is applied: • **Inbound** The policy is applied to traffic as it enters the interface. • **Outbound** The policy is applied to traffic as it exits the interface.
Status	The status of the policy on the interface. A policy is *Up* if DiffServ is globally enabled, and if the interface is administratively enabled and has a link. Otherwise, the status is *Down*.
Policy	The DiffServ policy associated with the interface.
Configure Service dialog box fields – When you click **Add** or **Edit**, this dialog box opens and allows you to configure DiffServ interface policies. Specifying *None* for a policy has no effect when adding or editing interface policies. To remove an interface policy mapping, click **Remove** on the parent page. The following information describes the fields in this window.	
Interface	Select an interface to associate with a policy.
Policy In	The menu lists all policies configured with a type of *In*. Select the policy to apply to traffic as it enters the interface.
Policy Out	The menu lists all policies configured with a type of *Out*. Select the policy to apply to traffic as it exits the interface.

Use the buttons to perform the following tasks:

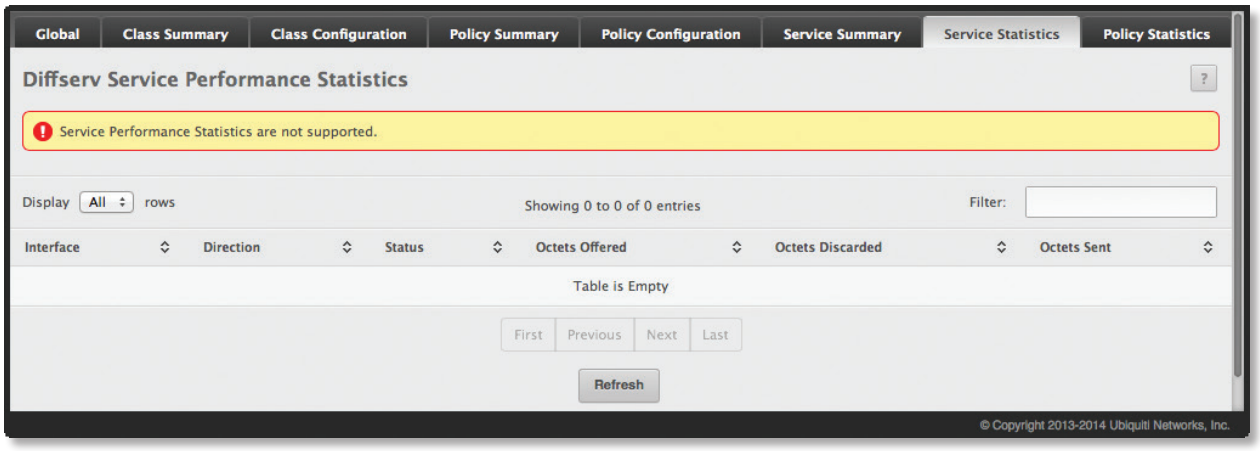
- To add a policy to an interface, click **Add**, configure the fields, and click **Submit** to apply the changes.
- To edit a configured interface-policy association, select the entry to modify and click **Edit**. Configure the fields, and click **Submit** to apply the changes.
- To remove one or more 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

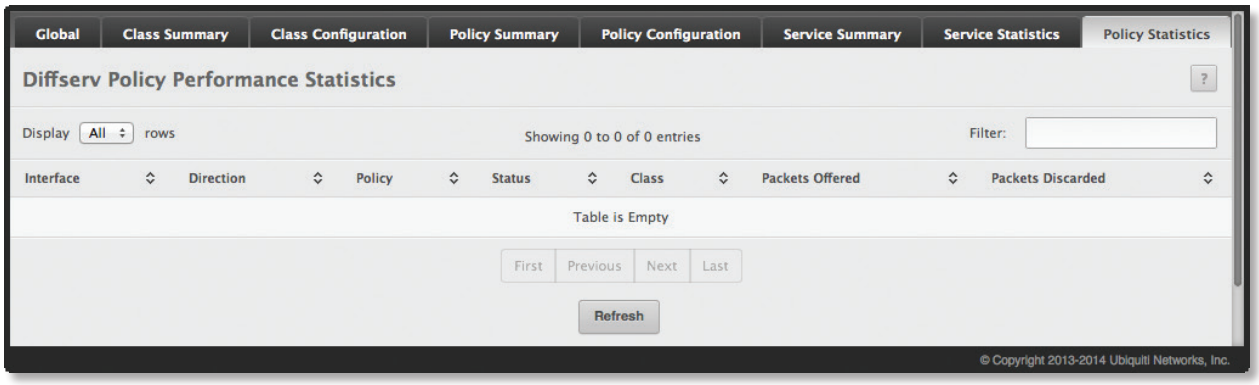
Field	Description
Interface	The interface associated with the rest of the data in the row. The table displays all interfaces that have a DiffServ policy currently attached in a traffic flow direction.
Direction	The traffic flow direction to which the policy is applied: • **In** The policy is applied to traffic as it enters the interface. • **Out** The policy is applied to traffic as it exits the interface.
Status	The operational status of this service interface, either *Up* or *Down*.
Octets Offered	The total number of octets offered to all class instances in this service policy before their defined DiffServ treatment is applied. This is the overall count per-interface, per-direction.
Octets Discarded	The total number of octets discarded for all class instances in this service policy for any reason due to DiffServ treatment. This is the overall count per-interface, per-direction.
Octets Sent	The total number of octets forwarded for all class instances in this service policy after their defined DiffServ treatments were applied. In this case, forwarding means the traffic stream was passed to the next functional element in the data path, such as the switching or routing function of an outbound link transmission element. This is the overall count per-interface, per-direction.

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

Diffserv 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
Interface	The interface associated with the rest of the data in the row. The table displays all interfaces that have a DiffServ policy currently attached in a traffic flow direction.
Direction	The traffic flow direction to which the policy is applied: • **In** The policy is applied to traffic as it enters the interface. • **Out** The policy is applied to traffic as it exits the interface.
Policy	The name of the policy currently attached to the interface.
Status	The operational status of the policy currently attached to the interface.
Class	The DiffServ class currently defined for the attached policy.
Packets Offered	The total number of packets offered to all class instances in this service policy before their defined DiffServ treatment is applied. This is the overall count per-interface, per-direction.
Packets Discarded	The total number of packets discarded for all class instances in this service policy for any reason due to DiffServ treatment. This is the overall count per-interface, per-direction.

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

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:

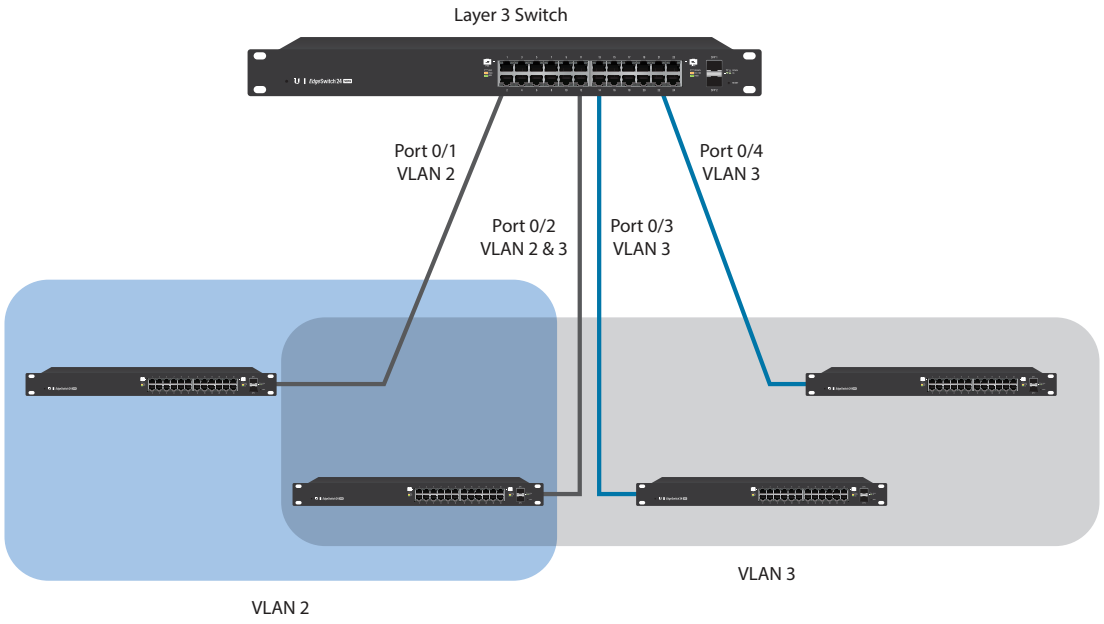
- **"Configuring VLANs" on page 259**
- **"Configuring Multiple Spanning Tree Protocol" on page 262**
- **"Configuring VLAN Routing" on page 265**
- **"Configuring Policy-Based Routing" on page 267**
- **"Configuring 802.1X Network Access Control" on page 270**
- **"Configuring Differentiated Services 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.

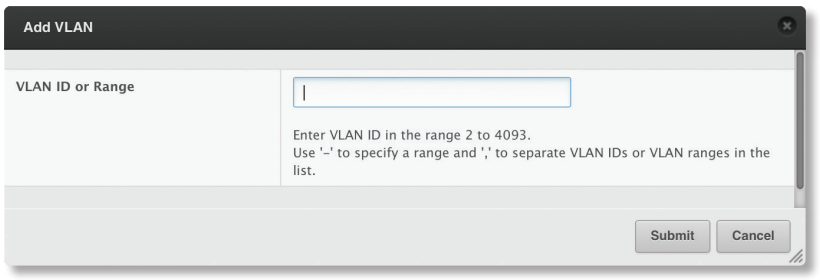


VLAN Example Network Diagram

Using the EdgeSwitch UI to Configure VLANs

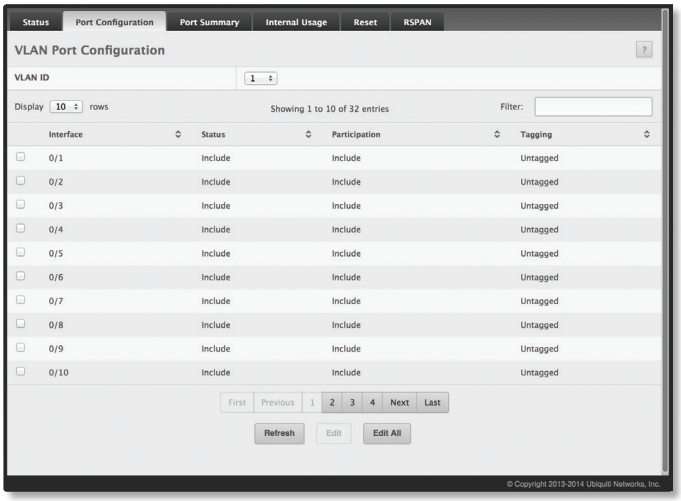
1. Access the **Switching** > **VLAN** > **Status** page.
2. Click **Add** to create a new VLAN.

3. Type **2-3** in the *VLAN ID or Range* field, and then click **Submit**.



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.



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

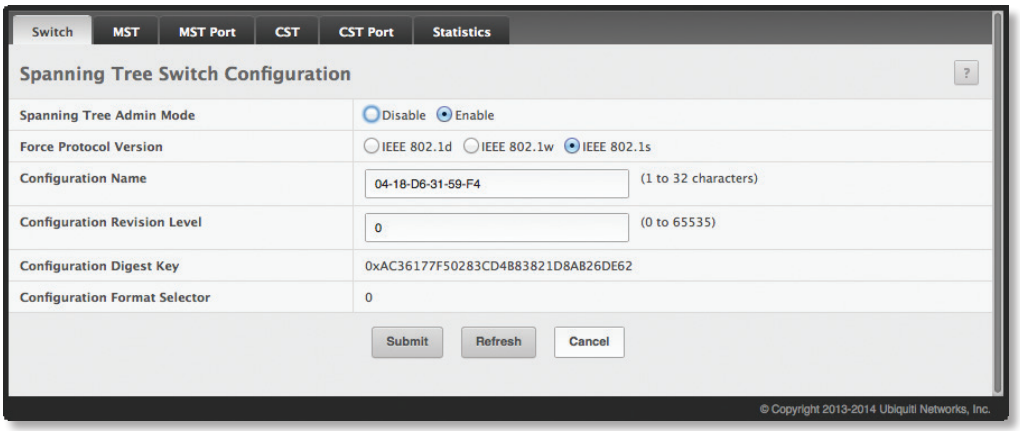
This example shows how to enable IEEE 802.1s Multiple Spanning Tree (MST) protocol on the switch and all of the ports and to set the bridge priority.

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

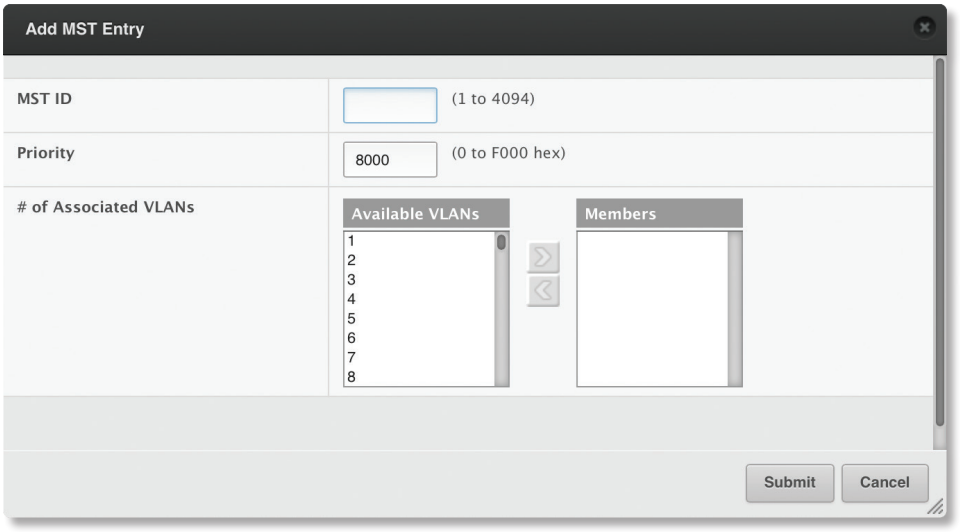
1. Create VLANs 10 and 20.
   a. Access the **Switching** > **VLAN** > **Status** page.
   b. Click **Add** to create a VLAN.
   c. Select the **VLAN ID-Individual** option and enter **10**.
   d. Click **Submit**.
   e. Repeat the steps to add VLAN 20.
2. Enable MSTP (IEEE 802.1s) on the switch and 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**.



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
```

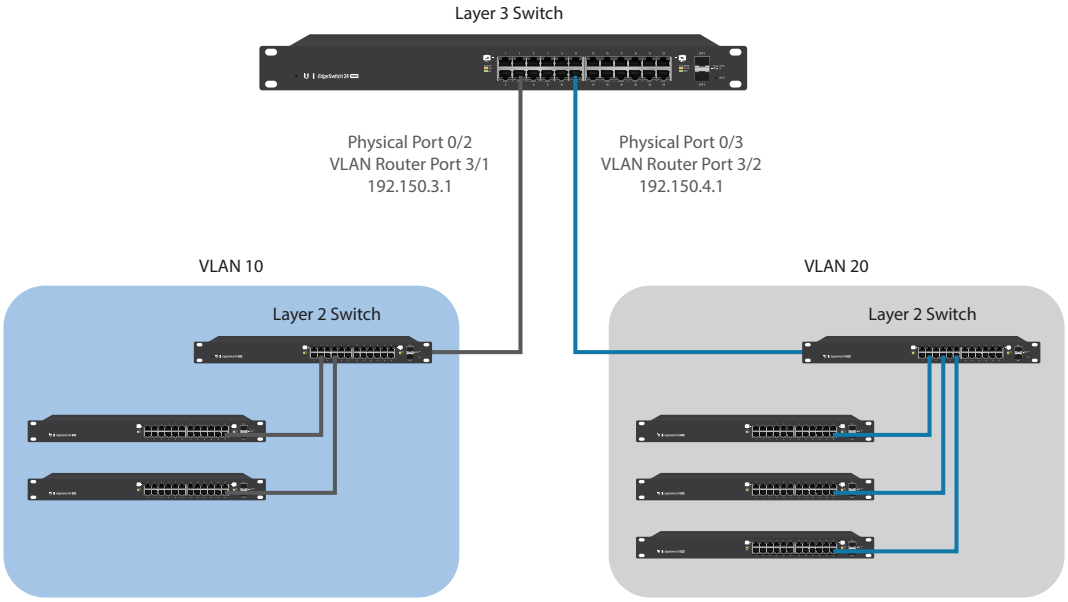
10. On the non-root bridge, change the priority to force port 0/2 to be the root port.

```
    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

           Logical
VLAN ID   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.

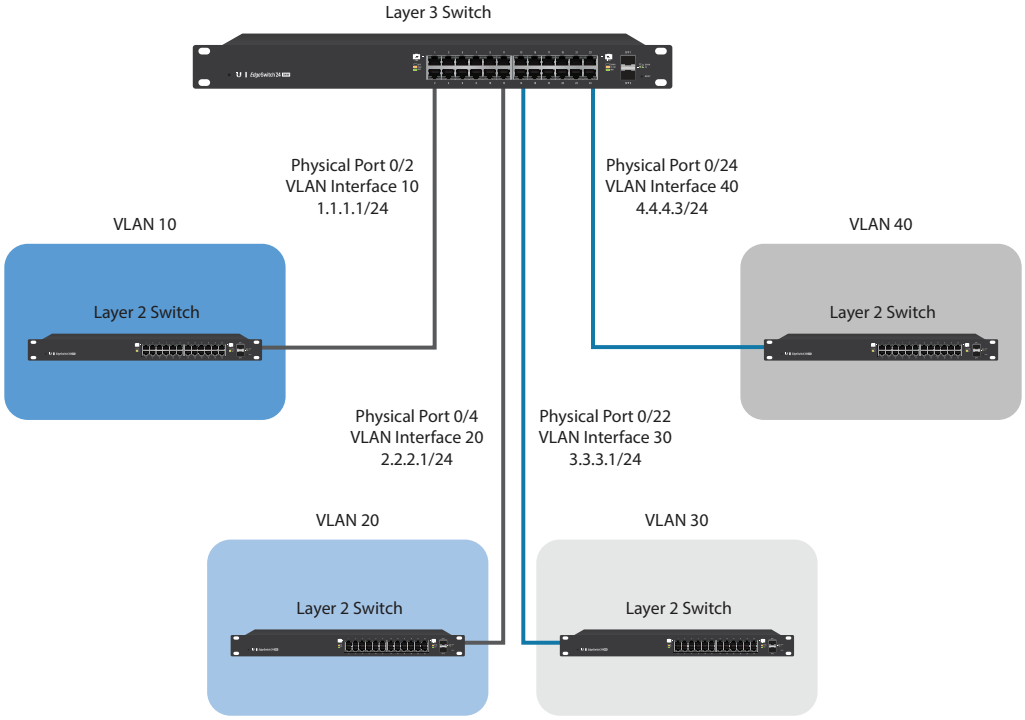
Configuring PBR involves configuring a route-map with `match` and `set` commands and then applying the corresponding route-map to the interface.

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
```

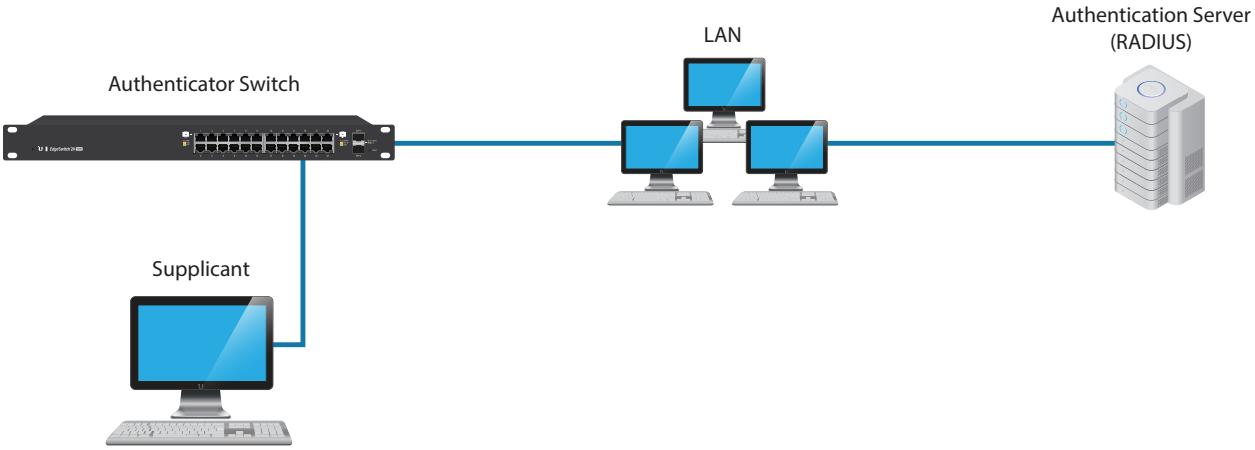
Set clauses:

```
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.



Switch with 802.1X Network Access Control

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.

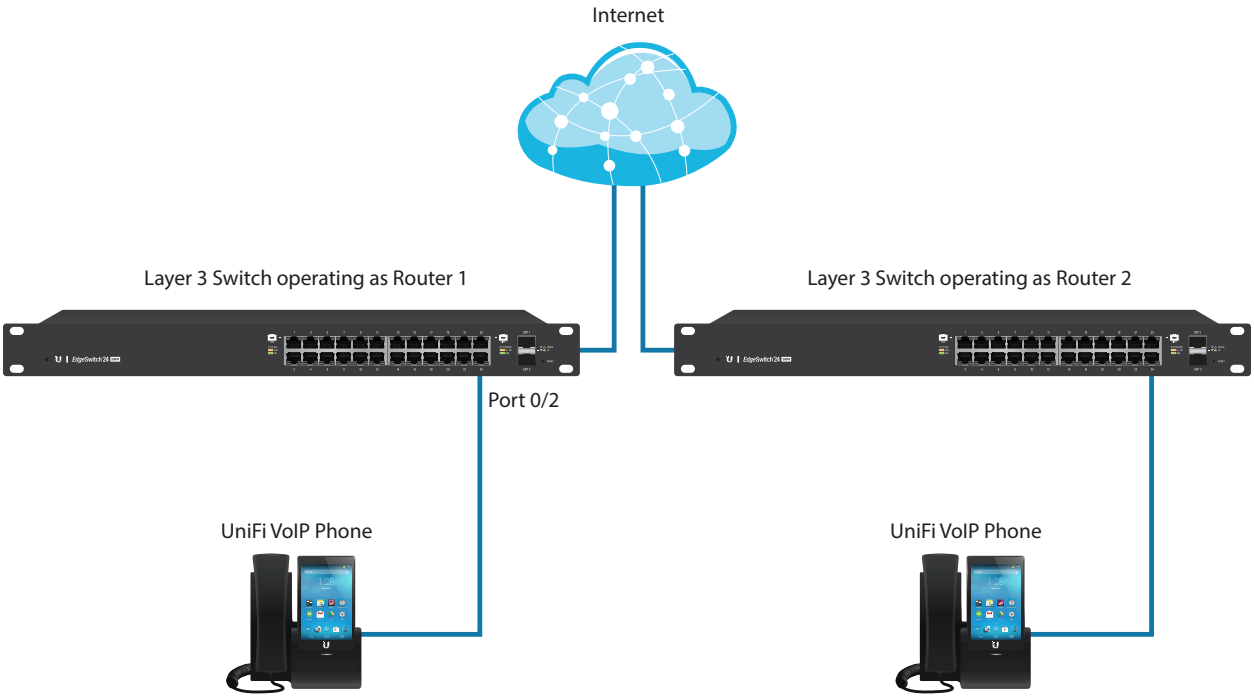
```
        dot1x system-auth-control
```

8. 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.



DiffServ VoIP Example Network Diagram

Using the CLI to Configure DiffServ VoIP Support

1. Enter Global Config mode. Set queue 5 on all ports to use strict priority mode. This queue shall be used for all VoIP packets. Activate DiffServ for the switch.

```
(UBNT EdgeSwitch) #config
    cos-queue strict 5
    diffserv
```

2. Create a DiffServ classifier named 'class_voip' and define a single match criterion to detect UDP packets. The class type match-all indicates that all match criteria defined for the class must be satisfied in order for a packet to be considered a match.

```
    class-map match-all class_voip
       match protocol udp
       exit
```

3. Create a second DiffServ classifier named ‘class_ef’ and define a single match criterion to detect a DiffServ code point (DSCP) of ‘EF’ (expedited forwarding). This handles incoming traffic that was previously marked as expedited elsewhere in the network.

```
class-map match-all class_ef
   match ip dscp ef
   exit
```

4. Create a DiffServ policy for inbound traffic named ‘pol_voip’, 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.

AI093014