kontron

»Kontron User's Guide«

FASTPATH
Administrator's Guide

Document Revision 1.0

Document ID: FASTPATH Administrator's Guide

Issue Date: November 2011

Revision History

Rev. Index	Brief Description of Changes	Date of Issue
1.0	Initial Issue	08.11.2011

Customer Service

Contact Information:

Kontron Canada, Inc.
4555 Ambroise-Lafortune
Boisbriand, Québec, Canada
J7H 0A4
Tel: (450) 437-5682
(800) 354-4223
Fax: (450) 437-8053
E-mail: support@ca.kontron.com

Kontron Modular Computer GmbH
Sudetenstrasse 7
87600 Kaufbeuren
Germany
+49 (0) 8341 803 333
+49 (0) 8341 803 339
support-kom@kontron.com

Visit our site at:www.kontron.com

© 2011 Kontron, an International Corporation. All rights reserved.

The information in this user's guide is provided for reference only. Kontron does not assume any liability arising out of the application or use of the information or products described herein. This user's guide may contain or reference information and products protected by copyrights or patents and does not convey any license under the patent rights of Kontron, nor the rights of others.

Kontron is a registered trademark of Kontron. All trademarks, registered trademarks, and trade names used in this user's guide are the property of their respective owners. All rights reserved. Printed in Canada. This user's guide contains information proprietary to Kontron. Customers may reprint and use this user's guide in other publications. Customers may alter this user's guide and publish it only after they remove the Kontron name, cover, and logo.

Kontron reserves the right to make changes without notice in product or component design as warranted by evolution in user needs or progress in engineering or manufacturing technology. Changes that affect the operation of the unit will be documented in the next revision of this user's guide.

Table of Contents

List of Tables

List of Figures

Proprietary Note

This document contains information proprietary to Kontron AG. It may not be copied or transmitted by any means, disclosed to others, or stored in any retrieval system or media without the prior written consent of Kontron AG or one of its authorized agents.

The information contained in this document is, to the best of our knowledge, entirely correct. However, Kontron AG cannot accept liability for any inaccuracies or the consequences thereof, or for any liability arising from the use or application of any circuit, product, or example shown in this document.

Kontron AG reserves the right to change, modify, or improve this document or the product described herein, as seen fit by Kontron AG without further notice.

Trademarks

Kontron AG and the *Kontron* logo are trade marks owned by Kontron AG, Germany. In addition, this document may include names, company logos and trademarks, which are registered trademarks and, therefore, proprietary to their respective owners.

About This Document

This guide describes how to configure the FASTPATH® software features by using the Web-based graphical user interface (GUI). The FASTPATH architecture accommodates a variety of software modules so that a platform running FASTPATH software can be a Layer 2 switch in a basic network or a Layer 3 router in a large, complex network.

Kontron Support for Switch Software

In case of support questions related to the Fastpath software on any of the products, please contact Kontron Support. Contact details are given in the corresponding product User's Guide.

To be able to process support cases as fast as possible, please add the following information:

- Output of
  - show boardinfo version
  - show tech-support
- Information of use-case
  - Overall system setup
  - Block diagram of used I/Fs and connected devices
  - Configuration of external devices (ETHx setup, ...)

Audience

The information in this guide is intended for any of the following individuals:

- System administrators who are responsible for configuring and operating a network using FASTPATH software
- Software engineers who are integrating FASTPATH software into a router or switch product
- Level 1 and/or Level 2 Support providers

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

Organization

This document is organized as follows:

- Section 1: "Getting Started" on page 1 contains information about performing the initial system configuration and accessing the user interfaces.
- Section 2: "Configuring System Information" on page 15 describes how to configure administrative features such as SNMP, DHCP, and port information.
- Section 3: "Configuring Switching Information" on page 150 describes how to manage and monitor the layer 2 switching features.
- Section 4: "Configuring Routing" on page 259 describes how to configure the layer 3 routing features.
- Section 5: "Managing Device Security" on page 330 contains information about configuring switch security information such as captive portal configuration, port access control, TACACS+, and RADIUS server settings.
- Section 6: "Configuring Quality of Service" on page 375 describes how to manage the FASTPATH software ACLs, and how to configure the Differentiated Services and Class of Service features.
- Section 7: "Configuring IP Multicast" on page 430 describes how to configure the IP multicast features.
- Appendix A:"Configuration Examples" on page 1 describe how to configure selected features on the switch by using the Web interface, command-line interface, and Simple Network Management Protocol (SNMP).

Additional Documentation

The following documentation provides additional information about FASTPATH software:

- The FASTPATH CLI Reference manual of the corresponding product describes the commands available from the command-line interface (CLI) for managing, monitoring, and configuring the switch.

Advisory Conventions

This section describes the conventions this document uses.

CAUTION

This symbol provides information about critical aspects of the configuration, combinations of settings, events or procedures that can adversely affect network connectivity, security and so on.

Note...

This symbol and title emphasize aspects the reader should read through carefully for his or her own advantage.

Typographical Conventions

This guide uses the typographical conventions described in the table below.

Symbol	Description	Example
Bold	Menu titles, button names and keyboard names when referred to in steps	Click **Submit** to apply your settings.
Blue Text	Hyperlinked text.	See "About This Document" on page xxvii.
`courier font`	Command-line text and file names	`(switch-prompt)#`

Fastpath Software Modules

The FASTPATH software suite includes the following modules:

- Switching (Layer 2)
- Routing (Layer 3)
- Multicast
- Quality of Service
- Management (CLI, Web UI and SNMP)

Not all modules are available for all platforms or software releases used on Kontronproducts.

FASTPATH software consists of flexible modules that can be applied in various combinations to develop advanced Layer 2/3/4+ products. The user-configurable features available on your switch depend on the installed module.

Note: Kontron products does not support stacking, thus the <Slot/Port> (e.g. 0/12) numbering is used instead of the <Unit/Slot/Port> (e.g. 1/0/12) is used. In case there are figures showing the <Unit/Slot/Port> numbering, the leading 1/ must be ignored.

Two Year Warranty

Kontron AG grants the original purchaser of Kontron's products a *TWO YEAR LIMITED HARDWARE WARRANTY* as described in the following. However, no other warranties that may be granted or implied by anyone on behalf of Kontron are valid unless the consumer has the express written consent of Kontron AG.

Kontron AG warrants their own products, excluding software, to be free from manufacturing and material defects for a period of 24 consecutive months from the date of purchase. This warranty is not transferable nor extendible to cover any other users or long-term storage of the product. It does not cover products which have been modified, altered or repaired by any other party than Kontron Modular Computers GmbH or their authorized agents. Furthermore, any product which has been, or is suspected of being damaged as a result of negligence, improper use, incorrect handling, servicing or maintenance, or which has been damaged as a result of excessive current/voltage or temperature, or which has had its serial number(s), any other markings or parts thereof altered, defaced or removed will also be excluded from this warranty.

If the customer's eligibility for warranty has not been voided, in the event of any claim, he may return the product at the earliest possible convenience to the original place of purchase, together with a copy of the original document of purchase, a full description of the application the product is used on and a description of the defect. Pack the product in such a way as to ensure safe transportation (see our safety instructions).

Kontron provides for repair or replacement of any part, assembly or sub-assembly at their own discretion, or to refund the original cost of purchase, if appropriate. In the event of repair, refunding or replacement of any part, the ownership of the removed or replaced parts reverts to Kontron Modular Computers GmbH, and the remaining part of the original guarantee, or any new guarantee to cover the repaired or replaced items, will be transferred to cover the new or repaired items. Any extensions to the original guarantee are considered gestures of goodwill, and will be defined in the "Repair Report" issued by Kontron with the repaired or replaced item.

Kontron Modular Computers GmbH will not accept liability for any further claims resulting directly or indirectly from any warranty claim, other than the above specified repair, replacement or refunding. In particular, all claims for damage to any system or process in which the product was employed, or any loss incurred as a result of the product not functioning at any given time, are excluded. The extent of Kontron Modular Computers GmbH liability to the customer shall not exceed the original purchase price of the item for which the claim exists.

Kontron Modular Computers GmbH issues no warranty or representation, either explicit or implicit, with respect to its products' reliability, fitness, quality, marketability or ability to fulfil any particular application or purpose. As a result, the products are sold "as is," and the responsibility to ensure their suitability for any given task remains that of the purchaser. In no event will Kontron be liable for direct, indirect or consequential damages resulting from the use of our hardware or software products, or documentation, even if Kontron were advised of the possibility of such claims prior to the purchase of the product or during any period since the date of its purchase.

Please remember that no Kontron Modular Computers GmbH employee, dealer or agent is authorized to make any modification or addition to the above specified terms, either verbally or in any other form, written or electronically transmitted, without the company's consent.

1 Getting Started

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

- Connecting the Switch to the Network
- Booting the Switch
- Understanding the User Interfaces

1.1 Connecting the Switch to the Network

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

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

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

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

Note...

Some switches provide a Service port, an Ethernet port usually located on the back on the switch, as a dedicated management port. On switches without a Service port, you use one of the network ports.

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

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

Follow these steps to make a serial port connection and configure network information:

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

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

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

2. Configure the terminal-emulation program to use the following settings:
   - Baud rate: 9600 bps
   - Data bits: 8
   - Parity: none
   - Stop bit: 1
   - Flow control: none
3. Power on the switch.

For information about the boot process, including how to access the boot menu, see 1.2Booting the Switch3.

4. Press the return key, and the `User:` prompt appears.

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

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

5. At the `(switch)>` prompt, enter `enable` to enter the Privileged EXEC command mode. There is no default password to enter Privileged EXEC mode. Press ENTER at the password prompt if you did not change the default password.

The command prompt changes to `(switch)#`.

6. Configure network information.

If the unit has a service port:

- To have the address assigned through DHCP:

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

- To use BootP, change the protocol by entering:
  `serviceport protocol bootp`
- To disable DHCP/BootP and manually assign an IPv4 address, enter:
  `serviceport protocol none`
  `serviceport ip <ipaddress> <netmask> [<gateway>]`, for example:
  `serviceport ip 192.168.2.23 255.255.255.0 192.168.2.1`
- To disable DHCP/BootP and manually assign an IPv6 address and (optionally) default gateway, enter:
  `serviceport protocol none`
  `serviceport ipv6 address <address>/<prefix-length> [eui64]`
  `serviceport ipv6 gateway <gateway>`

To view the assigned or configured network address, enter:
`show serviceport`

If the unit does not have a service port:

- To use a DHCP server to obtain the IP address, subnet mask, and default gateway information, enter:
  `network protocol dhcp`.
- To use a BootP server to obtain the IP address, subnet mask, and default gateway information, enter:
  `network protocol bootp`.
- To manually configure the IPv4 address, subnet mask, and (optionally) default gateway, enter:
  `network parms <ipaddress> <netmask> [<gateway>]`, for example:
  `network parms 192.168.2.23 255.255.255.0 192.168.2.1`
- To manually configure the IPv6 address, subnet mask, and (optionally) default gateway, enter:
  `network ipv6 address <address>/<prefix-length> [eui64]`
  `network ipv6 gateway <gateway>`

To view the network information, enter `show network`.

7. To save these changes so they are retained during a switch reset, enter the following command:

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

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

1.2 Booting the Switch

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

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

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

To boot the switch, perform the following steps:

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

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

4. During boot, you can use the Boot menu, if necessary, to run special procedures. To enter the Boot menu, press **2** within the first ten seconds after the following message appears.

```
Select an option. If no selection in 10 seconds then
operational code will start.

1 - Start operational code.
2 - Start Boot Menu.
Select (1, 2):2
```

For information about the Boot menu, see “1.2.1Boot Menu Functions3."

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

After the switch boots successfully, the User login prompt appears and you can use the local terminal to begin configuring the switch. However, before configuring the switch, make sure that the software version installed on the switch is the latest version. If it is not the latest version, download and install the latest version. See Download File To Switch (TFTP)127.

1.2.1 Boot Menu Functions

Note...

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

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

Use the following procedures to display the Boot menu:

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

```
Boot Menu Version: 12 jun 2007
Select an option. If no selection in 10 seconds then
operational code will start.

1 - Start operational code.
2 - Start Boot Menu.
Select (1, 2):2

Boot Menu Version: 12 jun 2007

Options available
1  - Start operational code
2  - Change baud rate
3  - Retrieve event log using XMODEM
4  - Load new operational code using XMODEM
5  - Load configuration using XMODEM
6  - Display operational code vital product data
7  - Run flash diagnostics
8  - Update boot code
9  - Delete operational code
10 - Reset the system
11 - Restore configuration to factory defaults (delete config files)
12 - Activate Backup Image
[Boot Menu]
```

The following sections describe the Boot menu options. If no selection is made within 10 seconds (default), the operational code starts.

1.2.1.1 Start Operational Code

Use option 1 to resume loading the operational code.

To relaunch the boot process from the Boot menu:

1. On the **Boot menu**, select **1** and press **<Enter>**.

   The following prompt displays:

```
Operational Code Date: Thu Jun 8 12:51:44 2006
Uncompressing.....
                        50%                     100%
||||||||||||||||||||||||||||||||||||||||||||||||||
   1 File: bootos.c                        Line: 462 Task: ffffe00 EC: 2863311530
(0xaaaaaaaa)
(0 d 0 hrs 0 min 13 sec)
Timebase: 24.750275 MHz, MEM: 99.001100 MHz, PCI: 33.000366 MHz, CPU: 198.002200 MHz
PCI device BCM5675_A0 attached as unit 0.
PCI device BCM5695_B0 attached as unit 1.
PCI device BCM5695_B0 attached as unit 2.
PCI device BCM5673_A1 attached as unit 3.
PCI device BCM5673_A1 attached as unit 4.
Adding BCM transport pointers
Configuring CPUTRANS TX
Configuring CPUTRANS RX
st_state(0) = 0x0
st_state(1) = 0x3
st_state(2) = 0x2
```

1.2.1.2 Change Baud Rate

Use option **2** to change the baud rate of the serial interface.

To change the baud rate from the Boot menu:

1. On the **Boot menu**, select **2** and press **<Enter>**.

   The following prompt displays:

```
[Boot Menu]2
Select baud rate:
1 - 1200
2 - 2400
3 - 4800
4 - 9600
5 - 19200
6 - 38400
7 - 57600
8 - 115200
0 - no change
```

> **Note...**
>
> The selected baud rate takes effect immediately.

2. The bootup process resumes.

1.2.1.3 Retrieve Event Log Using XMODEM

Use option 3 to retrieve the event log and download it to your ASCII terminal.

To retrieve the event log from the Boot menu:

1. On the **Boot menu**, select **3** and press **<Enter>**.

   The following prompt displays:

```
[Boot Menu] 3
Sending event log, start XMODEM receive.....
File asciilog.bin Ready to SEND in binary mode
Estimated File Size 169K, 1345 Sectors, 172032 Bytes
Estimated transmission time 3 minutes 20 seconds
Send several Control-X characters to cancel before transfer starts.
```

2. The bootup process resumes.

1.2.1.4 Load New Operational Code Using XMODEM

Use option 4 when a new software version must be downloaded to replace corrupted files, update, or upgrade the system software.

To download software from the Boot menu:

1. On the **Boot menu**, select **4** and press **<Enter>**.

   The following prompt displays:

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

2. When using HyperTerminal, click **Transfer** on the **HyperTerminal** menu bar.
3. From the **Transfer** menu, click **Send File**.

The **Send File** window displays.

4. Enter the file path for the file to be downloaded.
5. Make sure the protocol is defined as XMODEM.
6. Click **Send**.

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

After software downloads, the switch reboots automatically.

1.2.1.5 Load Configuration Using XMODEM

Use option 5 when a new configuration file must be downloaded to replace the saved system configuration file.

To download software from the Boot menu:

1. On the **Boot menu**, select **5** and press **<Enter>**.

The following prompt displays:

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

2. When using HyperTerminal, click **Transfer** on the **HyperTerminal** menu bar.
3. From the **Transfer** menu, click **Send File**.

The **Send File** window displays.

4. Enter the file path for the file to be downloaded.
5. Make sure the protocol is defined as XMODEM.
6. Click **Send**.

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

1.2.1.6 Display Operational Code Vital Product Data

Use option 6 to view boot image information.

To display boot image information from the Boot menu:

1. On the **Boot menu**, select **6** and press **<Enter>**.

The following prompt displays:

```
[Boot Menu] 6
The following image is in the Flash File System:
File Name.....................................image1
CRC...........................................0xb017 (45079)
Target Device.................................0x00508541
Size..........................................0x8ec50c (9356556)
Number of Components..........................2
Operational Code Size.........................0x7ec048 (8306760)
Operational Code Offset.......................0x74 (116)
Operational Code FLASH flag...................1
Operational Code CRC..........................0x9B4D
```

```
Boot Code Version.............................1
Boot Code Size................................0x100000 (1048576)
Boot Code Offset..............................0x7ec0bc (8306876)
Boot Code FLASH flag..........................0
Boot Code CRC.................................0x1CB8
VPD - rel 0 ver 31 maint_lvl 0
      Timestamp - Thu Jun  8 12:51:44 2006
      File - pc62xxr0v31.stk[Boot Menu]
```

2. The bootup process resumes.

1.2.1.7 Run Flash Diagnostics

Use option 7 to run flash diagnostics. User action is confirmed with a Y/N question before executing the command.

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

1. On the **Boot menu**, select **6** and press **<Enter>**.

   The following prompt displays:

```
[Boot Menu] 7
Do you wish to run flash diagnostics? (Boot code region will not be tested.) (y/n): y
Input number of diagnostic iterations -> 1
Testing 2 x 28F128J3 base: 0xfe000000
Iterations remaining = 1

Erasing sector 0
Verify sector 0 erased
Writing sector 0
Erasing sector 1
Verify sector 1 erased
Writing sector 1
Erasing sector 2
Verify sector 2 erased
Writing sector 2
Erasing sector 3
Verify sector 3 erased
Writing sector 3
Erasing sector 4
Verify sector 4 erased
Writing sector 4
Erasing sector 5
Verify sector 5 erased
Writing sector 5
Erasing sector 6
Verify sector 6 erased
Writing sector 6
```

Note...

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

```
Flash Diagnostics passed
[Boot Menu]
```

2. The bootup process resumes.

1.2.1.8 Update Boot Code

Use option 8 to update the boot code in the flash memory. This option is only valid after loading new boot code using Boot Menu option 4. User action is confirmed with a Y/N question before executing the command.

To download software from the Boot menu:

1. On the **Boot menu**, select **8** and press **<Enter>**.

   The following prompt displays:

```
Do you wish to update Boot Code? (y/n) y
Erasing Boot Flash.....Done.
Wrote 0x10000 bytes.
Wrote 0x20000 bytes.
Wrote 0x30000 bytes.
Wrote 0x40000 bytes.
Wrote 0x50000 bytes.
Wrote 0x60000 bytes.
Boot code updated
```

2. The bootup process resumes.

1.2.1.9 Delete Operational Code

Use option 9 to delete the active image from the flash memory. User action is confirmed with a Y/N question before executing the command.

To delete the backup image from the Boot menu:

1. On the **Boot menu**, select **8** and press **<Enter>**.

   The following prompt displays:

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

2. The bootup process resumes.

1.2.1.10 Reset the System

Use option 10 to clear all flash and reset the system to its default setting. User action is confirmed with a Y/N question before executing the command.

To reset the system from the Boot menu:

1. On the **Boot menu**, select **10** and press **<Enter>**.

   The following prompt displays:

```
[Boot Menu] 10
Are you SURE you want to reset the system? (y/n):y
Boot code......
SDRAM 256

Boot Menu Version: Oct 20 2004
Select an option. If no selection in 10 seconds then operational code will start.

1 - Start operational code.
2 - Start Boot Menu.

Select (1, 2):2
```

2. The bootup process resumes.

1.2.1.11 Restore Configuration To Factory Defaults (Delete Configuration Files)

Use option 11 to load using the system default configuration and to boot without using the current startup configuration. Selecting 11 from the Boot Menu restores system defaults. Boot Sequence can then be started by selecting 1 from the Boot Menu.

To download software from the Boot menu:

1. On the **Boot menu**, select **11** and press **<Enter>**.

   The following prompt displays:

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

2. The bootup process resumes.

1.2.1.12 Activate Backup Image

Use option 12 to activate the backup image. The active image becomes the backup when this option is selected.

To activate the backup image:

1. From the **Boot menu**, select **12** and press **<Enter>**.

   The following message displays:

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

2. The bootup process resumes.

1.3 Understanding the User Interfaces

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

- Web User Interface
- Command-Line Interface (CLI)
- Simple Network Management Protocol (SNMP)

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

This guide describes how to use the Web-based interface to manage and monitor the system. For information about how to manage and monitor the system by using the CLI, see the *FASTPATH CLI Command Reference* and the *FASTPATH Configuration Guide*.

Note...

The Web configuration and monitoring pages and CLI commands available for each platform depend on the FASTPATH software modules installed. For more information about the modules, see About FASTPATH Software Moduleslv.

1.3.1 Using the Web Interface

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

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

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

1. Open a Web browser and enter the IP address of the switch in the Web browser address field.
2. Type the user name and password into the fields on the login screen, and then click **Login**.

The user name and password are the same as those you use to log on to the command-line interface. By default, the user name is *admin*, and there is no password. Passwords are case sensitive.

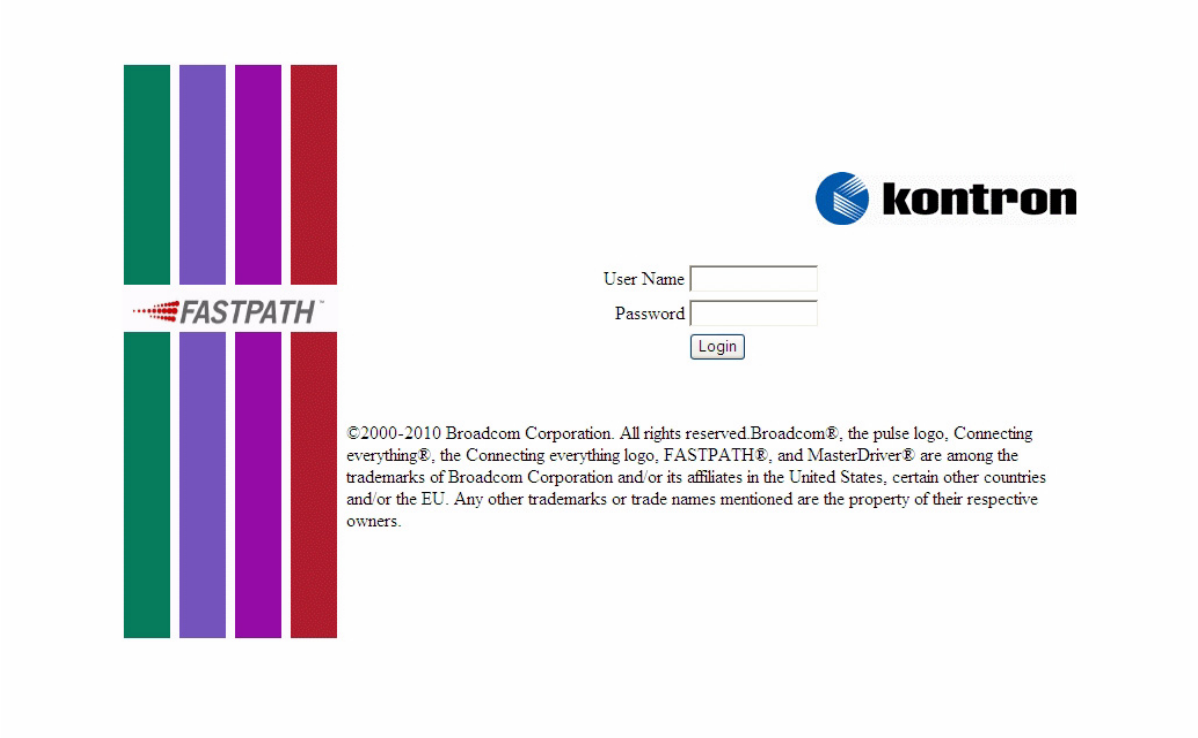

Figure 1-1: Login Page

3. After the system authenticates you, the System Description page displays.

Figure 1-2 shows the layout of the FASTPATH software Web interface. Each Web page contains two main areas: the navigation tree and the configuration status and options.

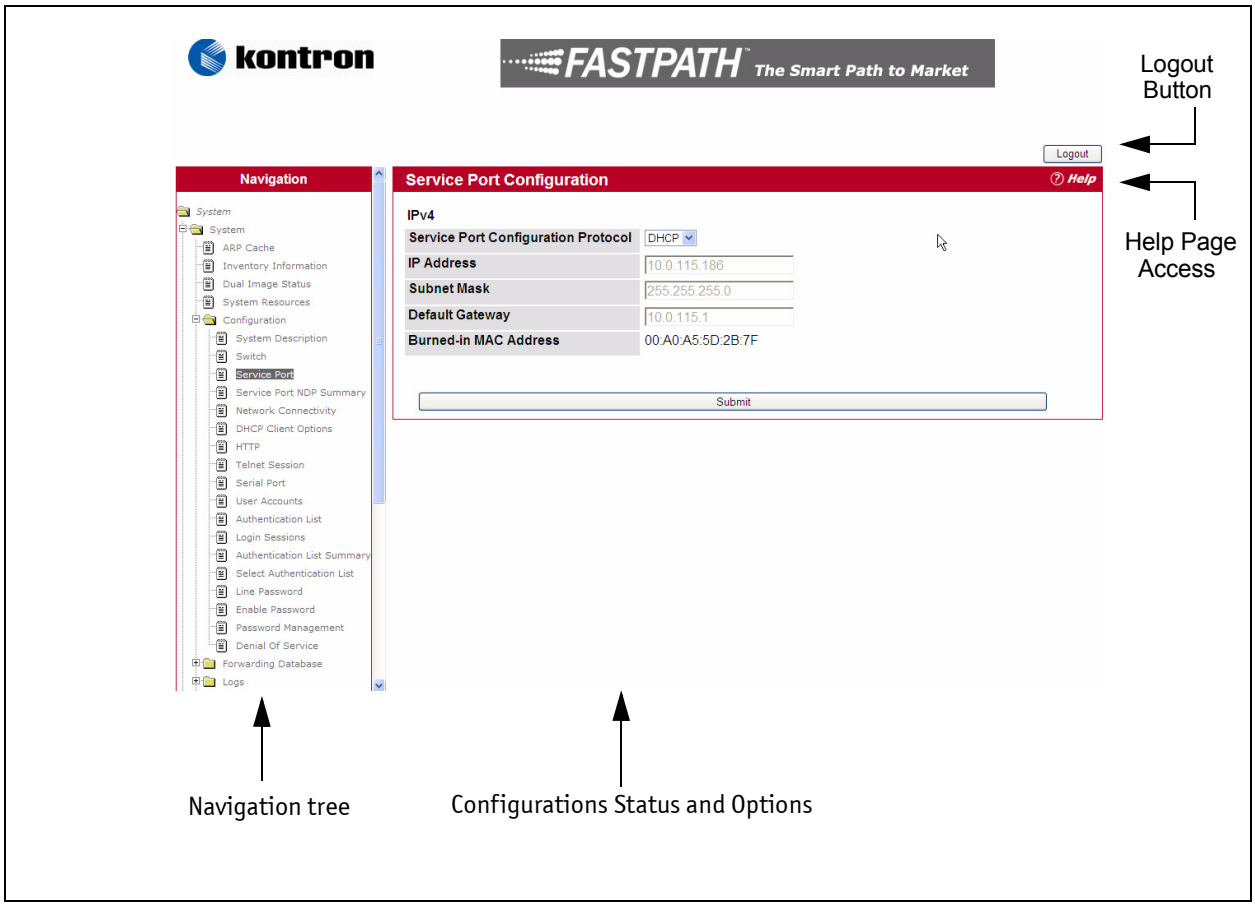

Figure 1-2: Web Interface Layout

1.3.1.1 Navigation Tree View

The hierarchical-tree view is on the left side of the Web interface. The tree view contains a list of various device features. The branches in the navigation tree can be expanded to view all the components under a specific feature, or retracted to hide the feature's components.

The tree consists of a combination of folders, subfolders, and configuration and status HTML pages. Click the folder to view the options in that folder. Each folder contains either subfolders or HTML pages, or a combination of both. Figure 1-3 shows an example of a folder, subfolder, and HTML page in the navigation menu. When you click a folder or subfolder that is preceded by a plus sign (+), the folder expands to display the contents. If you click an HTML page, a new page displays in the main frame. A folder or subfolder has no corresponding HTML page.

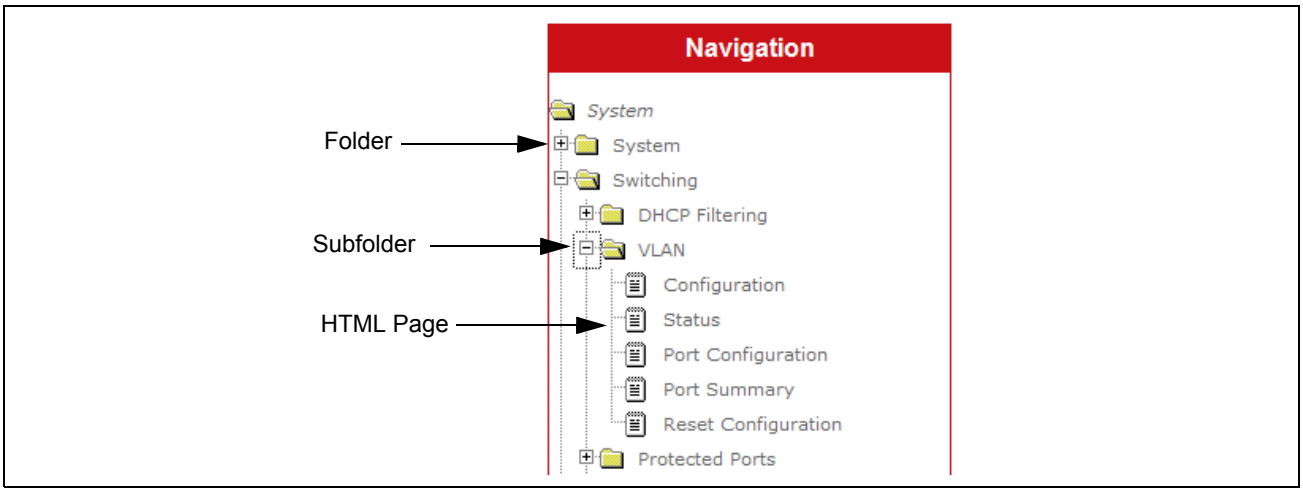

Figure 1-3: Navigation Tree View

1.3.1.2 Configuration and Monitoring Options

The panel to the right of the navigation menu displays the configuration information or status for the page you select. On pages that contain configuration options, you can input information into fields or select options from drop-down menus.

Each page contains access to the HTML-based help that explains the fields and configuration options for the page. Many pages also contain command buttons.

The following command buttons are used throughout the pages in the Web interface:

Table 1-1: Common Command Buttons

Button	Function
Submit	Clicking the **Submit** button sends the updated configuration to the switch. Configuration changes take effect immediately, but changes are not retained across a power cycle unless you save them to the system configuration file.
Refresh	Clicking the **Refresh** button refreshes the page with the latest information from the router.
Save	Clicking the **Save** button saves the current configuration to the system configuration file. When you click **Save**, changes that you have submitted are saved even when you reboot the system. To save the configuration to non-volatile memory, navigate to the **System > System Utilities > Save All Applied Changes** page and click **Save**.
Logout	Clicking the **Logout** button ends the session.

CAUTION

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

1.3.1.3 Help Page Access

Every page contains a link to the online help, which contains information to assist in configuring and managing the switch. The online help pages are context sensitive. For example, if the IP Addressing page is open, the help topic for that page displays if you click Help. Figure 1-4 shows the link to click to access online help on each page.

Figure 1-4: Help Link

Figure 1-2 on page 11 shows the location of the Help link on the Web interface.

1.3.1.4 User-Defined Fields

User-defined fields can contain 1-159 characters, unless otherwise noted on the configuration Web page.

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

```
\          <
/          >|
*          |
?
```

1.3.2 Using the Command-Line Interface

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

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

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

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

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

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

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

The `show` commands in the document also include a description of the information that the command shows.

1.3.3 Using SNMP

For FASTPATH software that includes the SNMP module, you can configure SNMP groups and users that can manage traps that the SNMP agent generates.

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

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

Any user can connect to the switch using the SNMPv3 protocol, but for authentication and encryption, you need to configure a new user profile. To configure a profile by using the CLI, see the SNMP section in the *FASTPATH CLI Command Reference*. To configure an SNMPv3 profile by using the Web interface, use the following steps:

1. Select **System > Configuration > User Accounts** from the hierarchical tree on the left side of the Web interface.
2. From the **User** menu, select **Create** to create a new user.
3. Enter a new user name in the **User Name** field.
4. Enter a new user password in the **Password** field and then retype it in the **Confirm Password** field.

To use SNMPv3 Authentication for this user, set a password of eight or more alphanumeric characters.

5. To enable authentication, use the **Authentication Protocol** menu to select either MD5 or SHA for the authentication protocol.
6. To enable encryption, use the **Encryption Protocol** menu to select **DES** for the encryption scheme. Then, enter an encryption code of eight or more alphanumeric characters in the Encryption Key field.
7. Click **Submit**.

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

2 Configuring System Information

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

- Viewing ARP Cache
- Viewing Inventory Information
- Viewing the Dual Image Status
- Viewing System Resources
- Defining General Device Information
- Configuring and Searching the Forwarding Database
- Managing Logs
- Configuring and Viewing Device Slot Information
- Configuring and Viewing Device Port Information
- TR-069 Client
- Configuring sFlow
- Defining SNMP Parameters
- Viewing System Statistics
- Using System Utilities
- Managing SNMP Traps
- Managing the DHCP Server
- Configuring DNS
- Configuring SNTP Settings
- Configuring and Viewing ISDP Information

2.1 Viewing ARP Cache

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

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

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

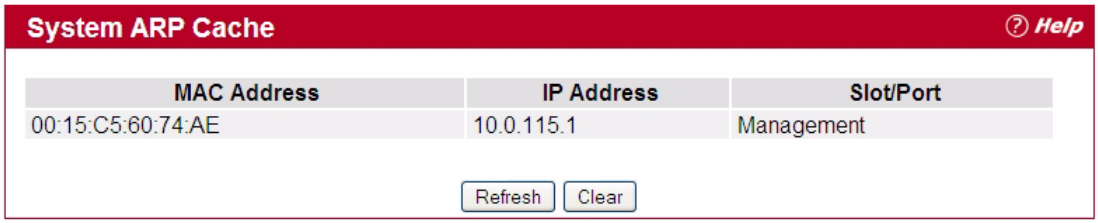

Figure 2-1: ARP Cache

Table 2-1: ARP Cache Fields

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

Click **Refresh** to reload the page and refresh the ARP cache view.

2.2 Viewing Inventory Information

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

To display the inventory information, click **System** > **Inventory Information** page in the navigation tree.

Inventory Information Help

Management Unit Number	1
System Description	Broadcom FASTPATH Routing
Machine Type	Broadcom HELIX 56304 Development System - 24 FE, 4 TENGIG
Machine Model	BCM-56304
Serial Number	n
FRU Number	
Part Number	BCM956304
Maintenance Level	A
Manufacturer	0xbc00
Base MAC Address	00:06:29:32:81:40
Software Version	L.4.3.1
Operating System	VxWorks5.5.1
Network Processing Device	BCM56304 REV 1
Additional Packages	FASTPATH BGP-4 FASTPATH QoS FASTPATH Multicast FASTPATH IPv6 FASTPATH Stacking

Refresh

Figure 2-2: Inventory Information

Table 2-2: Inventory Information Fields

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

2.3 Viewing the Dual Image Status

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

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

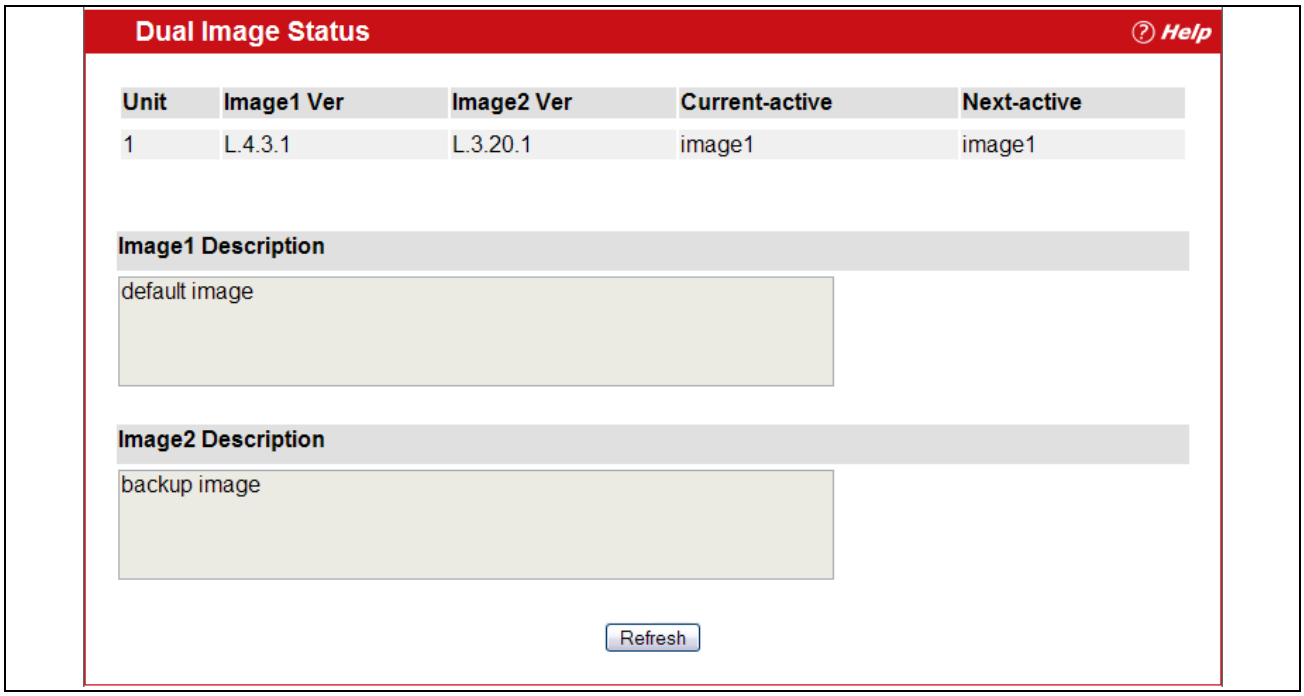

Figure 2-3: Dual Image Status

Table 2-3: Dual Image Status Fields

Field	Description
Unit	Displays the unit ID of the switch.
Image1 Ver	Displays the version of the image1 code file.
Image2 Ver	Displays the version of the image2 code file.
Current-active	Displays the currently active image on this unit.
Next-active	Displays the image to be used on the next restart of this unit.
Image1 Description	Displays the description associated with the image1 code file.
Image2 Description	Displays the description associated with the image2 code file.

- Click **Refresh** to display the latest information from the router.
- For information about how to update or change the system images, see 2.14Using System Utilities102.

2.4 Viewing System Resources

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

- Free memory
- Allocated memory
- CPU utilization by task

- Total CPU utilization at the following intervals:
  - Five seconds
  - One minute
  - Five minutes

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

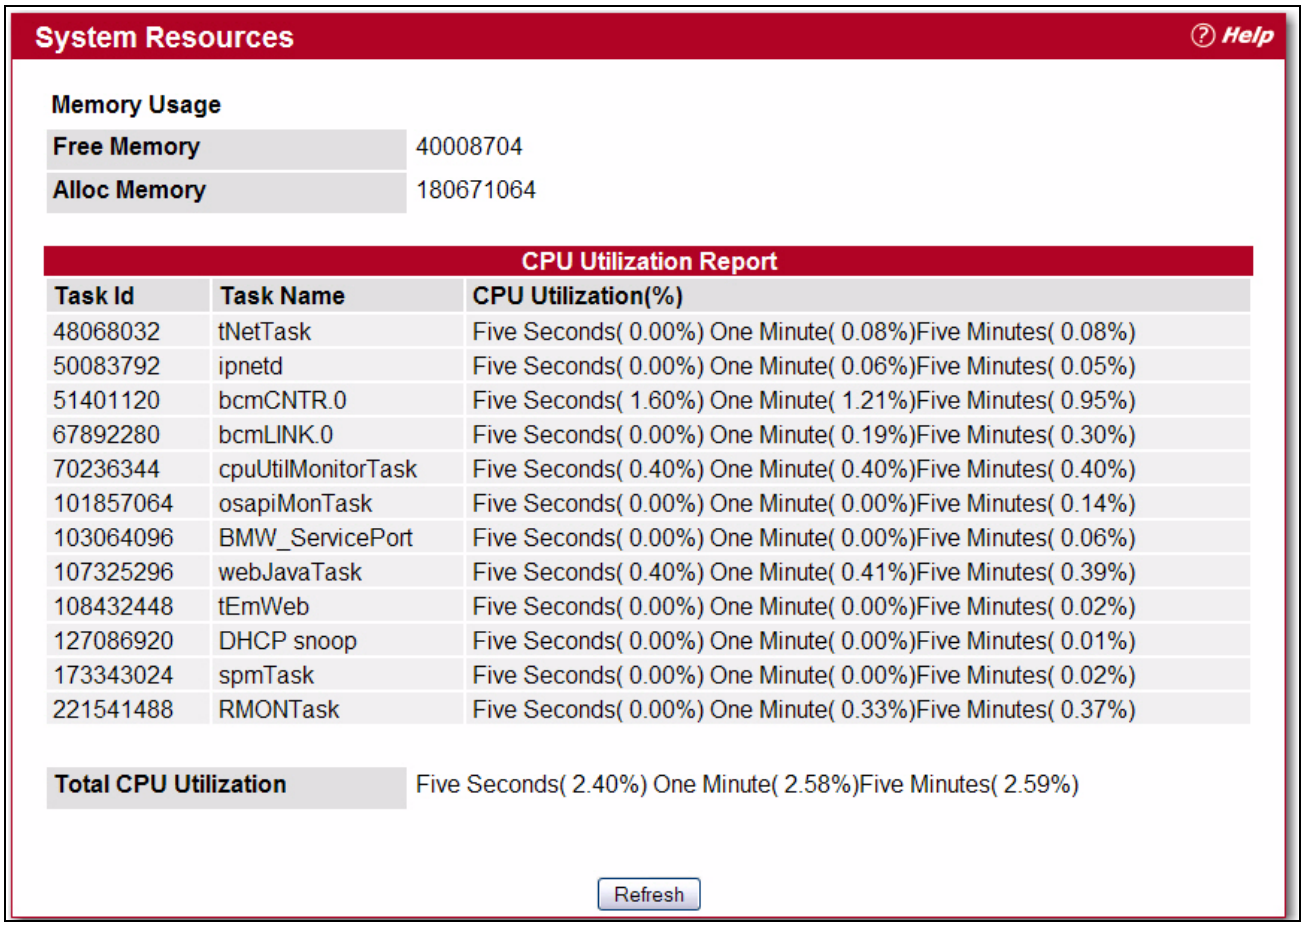

System Resources Help

Memory Usage

Free Memory	40008704
Alloc Memory	180671064

CPU Utilization Report

Task Id	Task Name	CPU Utilization(%)
48068032	tNetTask	Five Seconds(0.00%) One Minute(0.08%)Five Minutes(0.08%)
50083792	ipnetd	Five Seconds(0.00%) One Minute(0.06%)Five Minutes(0.05%)
51401120	bcmCNTR.0	Five Seconds(1.60%) One Minute(1.21%)Five Minutes(0.95%)
67892280	bcmLINK.0	Five Seconds(0.00%) One Minute(0.19%)Five Minutes(0.30%)
70236344	cpuUtilMonitorTask	Five Seconds(0.40%) One Minute(0.40%)Five Minutes(0.40%)
101857064	osapiMonTask	Five Seconds(0.00%) One Minute(0.00%)Five Minutes(0.14%)
103064096	BMW_ServicePort	Five Seconds(0.00%) One Minute(0.00%)Five Minutes(0.06%)
107325296	webJavaTask	Five Seconds(0.40%) One Minute(0.41%)Five Minutes(0.39%)
108432448	tEmWeb	Five Seconds(0.00%) One Minute(0.00%)Five Minutes(0.02%)
127086920	DHCP snoop	Five Seconds(0.00%) One Minute(0.00%)Five Minutes(0.01%)
173343024	spmTask	Five Seconds(0.00%) One Minute(0.00%)Five Minutes(0.02%)
221541488	RMONTask	Five Seconds(0.00%) One Minute(0.33%)Five Minutes(0.37%)

Total CPU Utilization	Five Seconds(2.40%) One Minute(2.58%)Five Minutes(2.59%)

Refresh

Figure 2-4: System Resources

Table 2-4: Dual Image Status Fields

Field	Description
Free Memory	Displays the available Free Memory on the switch.
Alloc Memory	Displays the allocated Memory for the switch.
Task Id	Displays the Id of running tasks.
Task Name	Displays the name of the running tasks.
CPU Utilization(%)	Displays the CPU Utilization of tasks in terms of percentage of utilization.
Total CPU Utilization	Displays the Total CPU Utilization in terms of percentage. **Total CPU Utilization** is shown in the following intervals: • Five seconds • One minute • Five minutes

2.5 Defining General Device Information

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

- System Description
- Switch ConfigurationService Port
- Service Port NDP Summary
- Service Port DHCPv6 Client Statistics
- Network Connectivity
- Network Connection NDP Summary
- Network Port DHCPv6 Client Statistics
- DHCP Client Options
- HTTP Configuration
- Telnet Session
- Serial Port
- User Accounts
- Authentication List Configuration
- Login Session
- Authentication List Summary
- Select Authentication List
- Line Password
- Enable Password
- Password Management
- Denial of Service

2.5.1 System Description

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

To display the System Description page, click **System > Configuration > System Description** in the navigation tree.

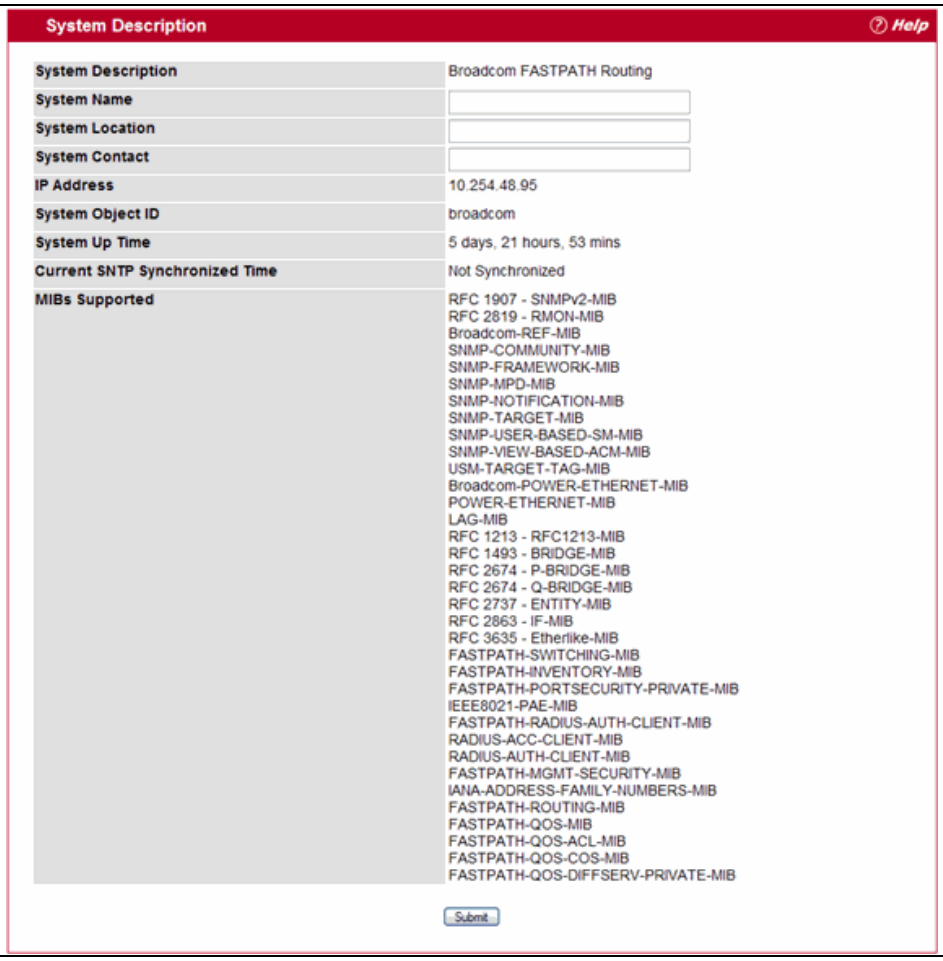

Figure 2-5: System Description

Table 2-5: System Description Fields

Field	Description
System Description	The product name of this switch.
System Name	Enter the name you want to use to identify this switch. You may use up to 31 alpha-numeric characters. The factory default is blank.
System Location	Enter the location of this switch. You may use up to 31 alpha-numeric characters. The factory default is blank.
System Contact	Enter the contact person for this switch. You may use up to 31 alpha-numeric characters. The factory default is blank.
IP Address	The IP Address assigned to the network interface. To change the IP address, see 2.5.6 Network Connectivity28.
System Object ID	The base object ID for the switch's enterprise MIB.
System Up Time	Displays the number of days, hours, and minutes since the last system restart.
Current SNTP Synchronized Time	Displays currently synchronized SNTP time in UTC. If no SNTP server has been configured and the time is not synchronized, this field displays “Not Synchronized.” To specify an SNTP server, see 2.18Configuring SNTP Settings138.
MIBs Supported	Displays the list of MIBs supported by the management agent running on this switch.

2.5.1.1 Defining System Information

1. Open the **System Description** page.
2. Define the following fields: **System Name**, **System Contact**, and **System Location**.
3. Click **Submit**.

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

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

2.5.2 Switch Configuration

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

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

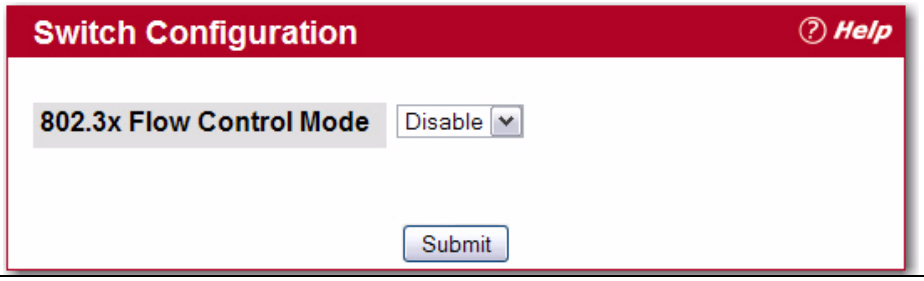

Figure 2-6: Switch 802.3x Flow Control

Table 2-6: Switch Configuration Fields

Field	Description
IEEE 802.3x Flow Control Mode	Enables or disables IEEE 802.3x flow control on the system. The factory default is disabled.
Enable	Enables flow control so that the switch can communicate with higher speed switches.
Disable	Disables flow control so that the switch does not send pause packets if the port buffers become full.

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

2.5.3 Service Port

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

To access the Service Port Configuration page, click **System > Configuration > Service Port** in the navigation tree.

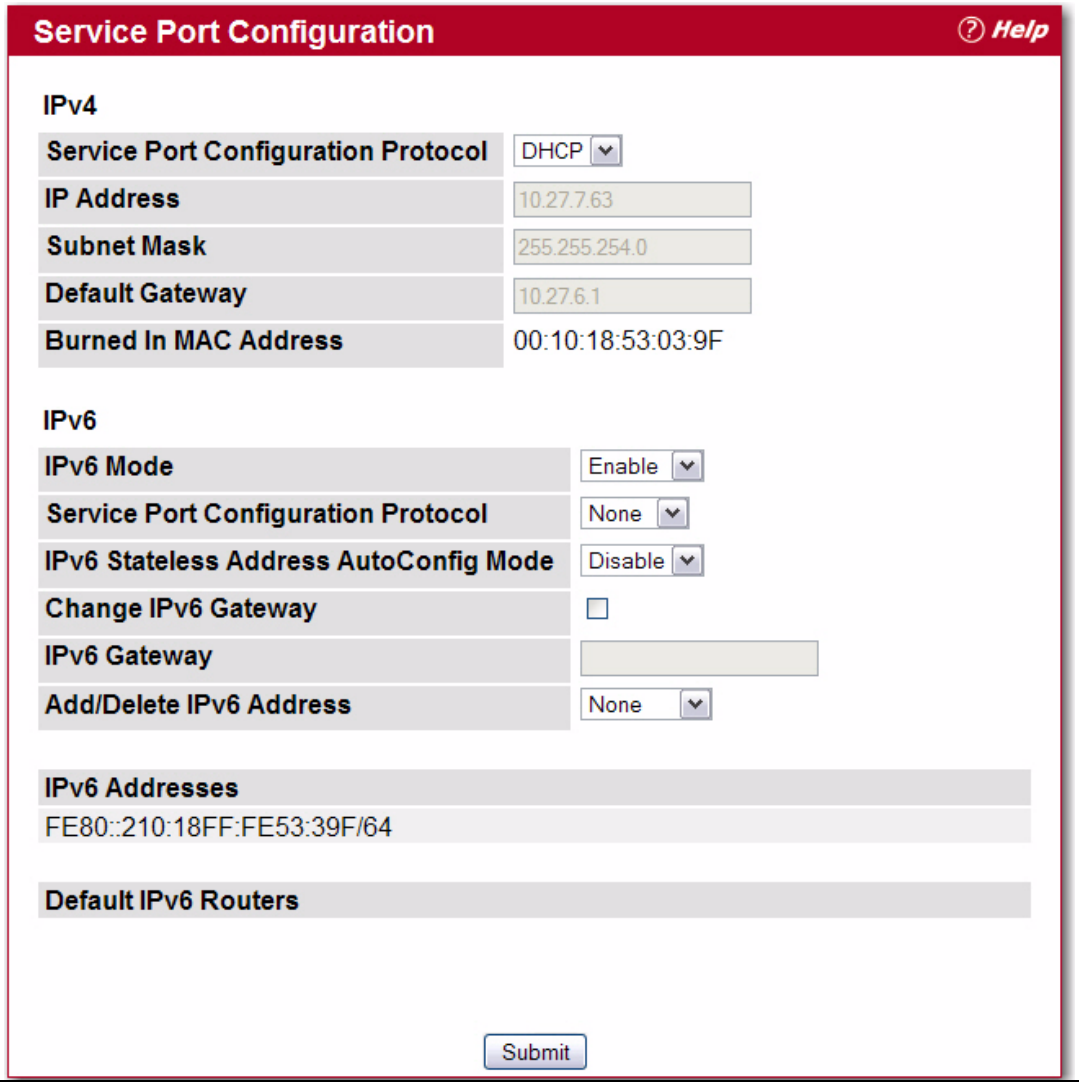

Figure 2-7: Service Port Configuration

Table 2-7: Service Port Configuration Fields

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

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

2.5.4 Service Port NDP Summary

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

To display the page, click **System > Configuration > Service Port NDP Summary**.

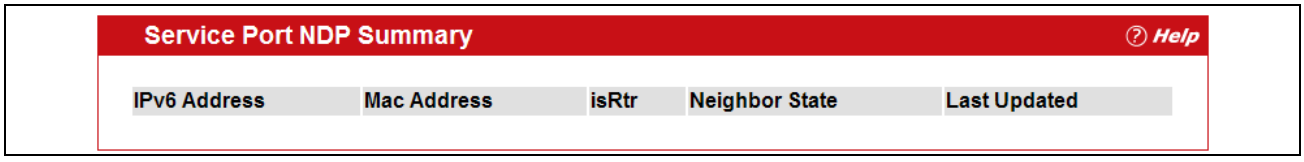

Figure 2-8: Service Port NDP Summary

Table 2-8: Service Port NDP Summary Fields

Field	Description
IPv6 Address	Displays the IP address of the neighbor.
MAC Address	Displays the MAC address of the neighbor.
isRtr	Indicates whether the neighbor is a router. If the neighbor is a router, the value is TRUE. If the neighbor is not a router, the value is FALSE.
Neighbor State	Specifies the state of the neighbor cache entry. Following are the states for dynamic entries in the IPv6 neighbor discovery cache: • Icmp: Address resolution is being performed on the entry. A neighbor solicitation message has been sent to the solicited-node multicast address of the target, but the corresponding neighbor advertisement message has not yet been received. • Reachable: Positive confirmation was received within the last Reachable Time milliseconds that the forward path to the neighbor was functioning properly. While in REACH state, the device takes no special action as packets are sent. • Stale: More than ReachableTime milliseconds have elapsed since the last positive confirmation was received that the forward path was functioning properly. While in STALE state, the device takes no action until a packet is sent. • Delay: More than ReachableTime milliseconds have elapsed since the last positive confirmation was received that the forward path was functioning properly. A packet was sent within the last DELAY_FIRST_PROBE_TIME seconds. If no reachability confirmation is received within DELAY_FIRST_PROBE_TIME seconds of entering the DELAY state, send a neighbor solicitation message and change the state to PROBE. • Probe: A reachability confirmation is actively sought by resending neighbor solicitation messages every RetransTimer milliseconds until a reachability confirmation is received.
Last Updated	Displays the time since the address was confirmed to be reachable.

2.5.5 Service Port DHCPv6 Client Statistics

The Service Port DHCPv6 Client Statistics page displays DHCPv6 client statistics.

To display the Service Port DHCP Client Statistics page, click **System > Configuration > Service Port DHCPv6 Client Statistics**.

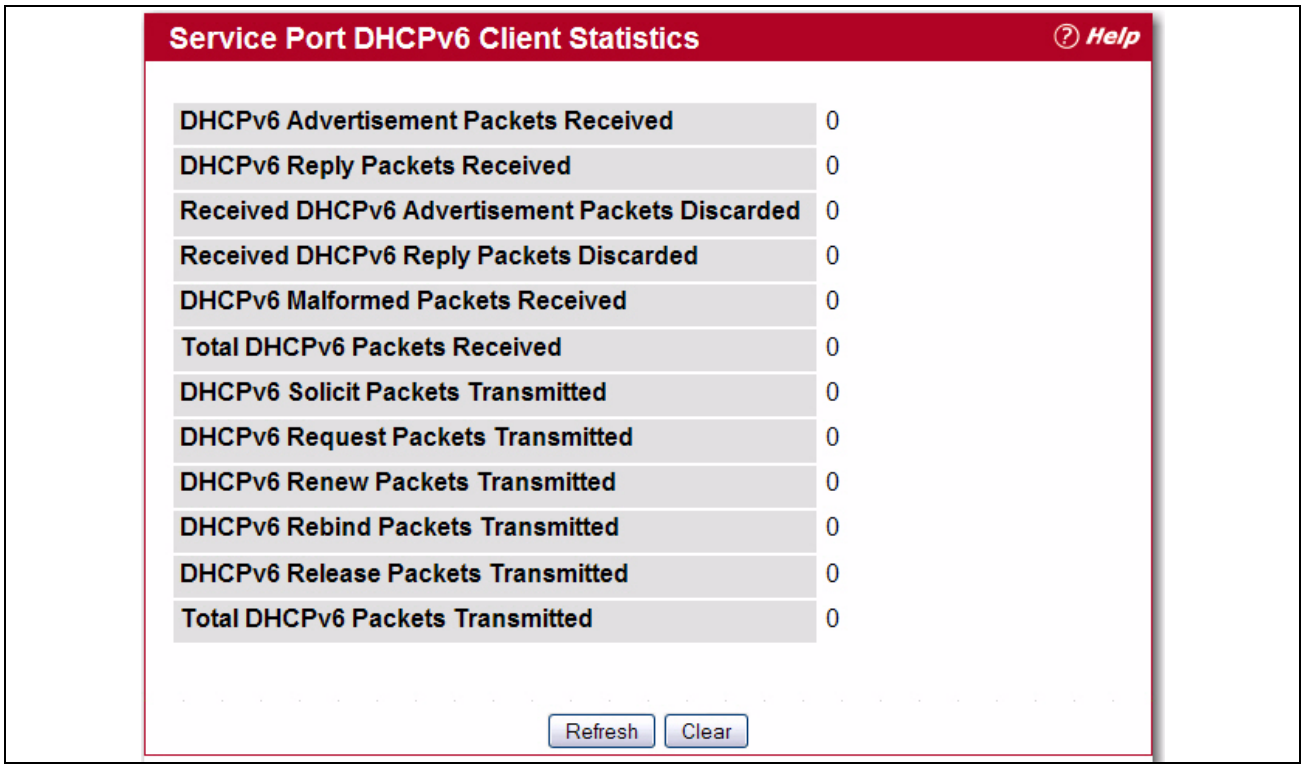

Figure 2-9: Service Port DHCPv6 Client Statistics

Table 2-9: Service Port DHCPv6 Client Statistics Fields

Field	Description
DHCPv6 Advertisement Packets Received	Displays the number of DHCPv6 Advertisement packets received on the service port.
DHCPv6 Reply Packets Received	Displays the number of DHCPv6 Reply packets received on the service port.
Received DHCPv6 Advertisement Packets Discarded	Displays the number of DHCPv6 Advertisement packets discarded on the service port.
Received DHCPv6 Reply Packets Discarded	Displays the number of DHCPv6 Reply packets discarded on the service port.
DHCPv6 Malformed Packets Received	Displays the Number of DHCPv6 packets that are received malformed on the service port.
Total DHCPv6 Packets Received	Displays the total number of DHCPv6 packets received on the service port.
DHCPv6 Solicit Packets Transmitted	Displays the number of DHCPv6 Solicit packets transmitted on the service port.
DHCPv6 Request Packets Transmitted	Displays the number of DHCPv6 Request packets transmitted on the service port.
DHCPv6 Renew Packets Transmitted	Displays the number of DHCPv6 Renew packets transmitted on the service port.
DHCPv6 Rebind Packets Transmitted	Displays the number of DHCPv6 Rebind packets transmitted on the service port.
DHCPv6 Release Packets Transmitted	Displays the number of DHCPv6 Release packets transmitted on the service port.
Total DHCPv6 Packets Transmitted	Displays the total number of DHCPv6 packets transmitted on the service port.

2.5.6 Network Connectivity

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

The Network Connectivity page allows you to change the IP information using the Web interface.

To access the page, click **System > Configuration > Network Connectivity** in the navigation tree.

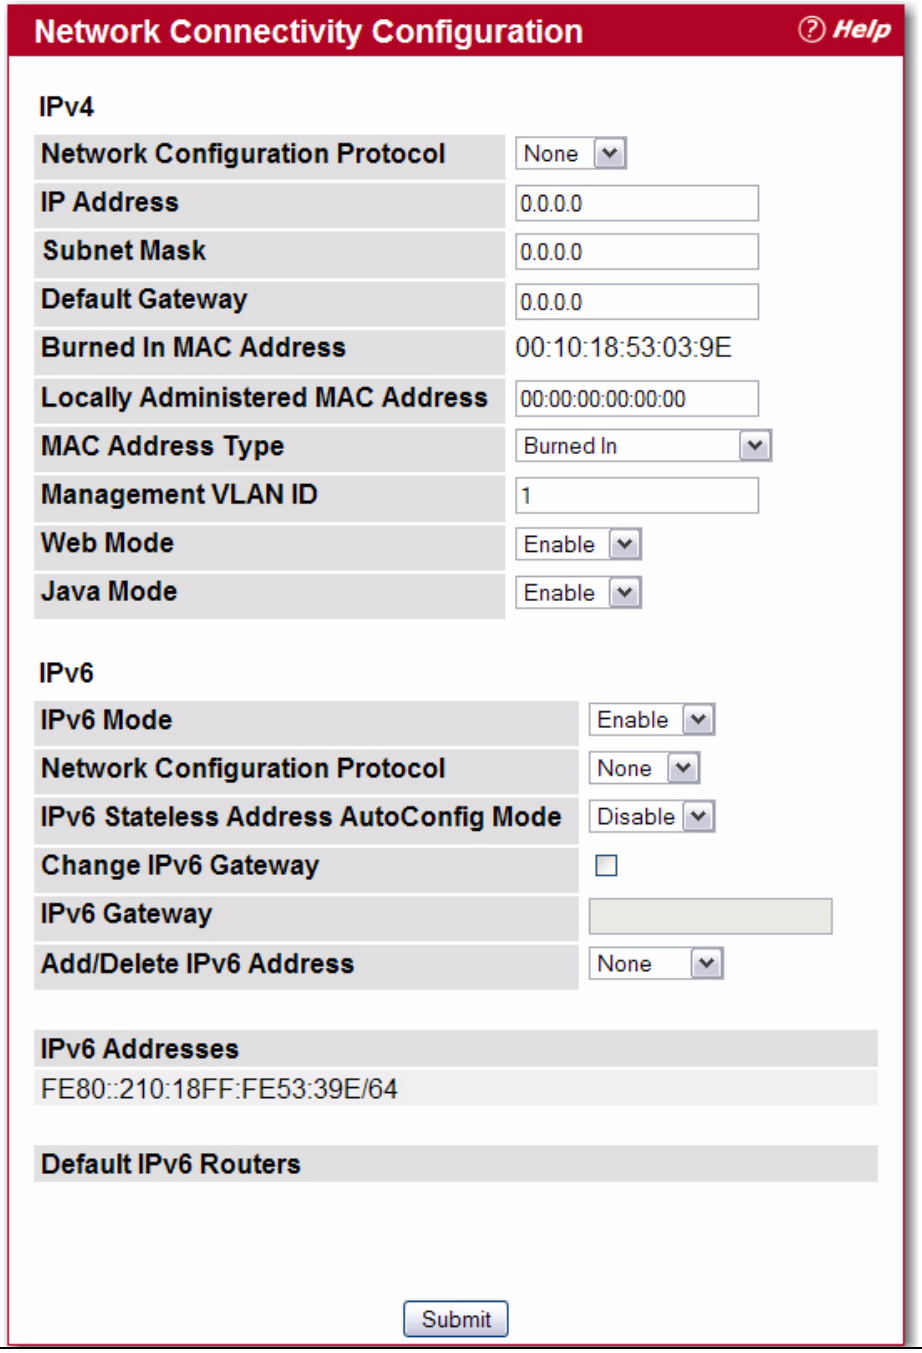

Figure 2-10: Network Connectivity Configuration

Table 2-10: Network Connectivity Configuration Fields

Field	Description
IPv4 Fields: These display IPv4 configuration information.	
Network Configuration Protocol	Specify what the switch should do following power-up. The factory default is None. The options are as follows: • **BootP**: Transmit a Bootp request. • **DHCP**: Transmit a DHCP request. • **None**: Do not send any requests following power-up.
IP Address	The IP address of the network interface. The factory default value is 0.0.0.0 ***Note:*** Each part of the IP address must start with a number other than zero. For example, IP addresses 001.100.192.6 and 192.001.10.3 are not valid.
Subnet Mask	The IP subnet mask for the interface. The factory default value is 0.0.0.0.
Default Gateway	The default gateway for the IP interface. The factory default value is 0.0.0.0.
Burned-in MAC Address	This read-only field displays the MAC address that is burned-in to the network card at the factory. This MAC address is used for in-band connectivity if you choose not to configure a locally administered address.
Locally Administered MAC Address	Specifies a locally administered MAC address for in-band connectivity instead of using the burned-in universally administered MAC address. In addition to entering an address in this field, you must also set the MAC address type to locally administered. Enter the address as twelve hexadecimal digits (6 bytes) with a colon between each byte. Bit 1 of byte 0 must be set to a 1 and bit 0 to a 0, i.e. byte 0 must have a value between x'40' and x'7F'.
MAC Address Type	Specify whether the burned-in or the locally administered MAC address should be used for in-band connectivity. The factory default is to use the burned-in MAC address
Management VLAN ID	Specify the management VLAN ID of the switch. It may be configured to any value in the range of (1 to 4093). The management VLAN is used for management of the switch. This field is configurable for administrative users and read-only for other users.
Web Mode	Enables/Disables Web Mode on the switch.
Java Mode	Enables/Disables Java mode on the switch.
IPv6 Fields: These display IPv6 configuration information.	
IPv6 Mode	Enable or disable IPv6 mode.
Network Configuration Protocol	Enable or Disable DHCPv6 Client protocol on the management port. The factory default is **None**.
IPv6 Stateless Address AutoConfig Mode	Enables or Disables the IPv6 stateless address autoconfiguration on the management port. The factory default is **None**.
DHCPv6 Client DUID	Displays the Client Identifier used by DHCPv6 Client when sending messages to the DHCPv6 Server. This entry displays only if IPv6 Network Configuration Protocol is set to DHCP.
Change IPv6 Gateway	Select the checkbox to configure an IPv6 Address.
IPv6 Gateway	Enter the IPv6 gateway address (do not include a prefix).
Add/Delete IPv6 Address	Select to **Add** or **Remove** IPv6 Addresses. The fields **New IPv6 Address** and **EUI Flag** are visible when we select the **Add** from this menu.
New IPv6 Address	Displays when **Add IPv6 Address** is selected. Adds IPv6 address.
EUI Flag	Displays when **Add IPv6 Address** is selected. Sets the EUI flag while configuring a new IPv6 address when **TRUE** is selected. The Default is **FALSE**.
IPv6 Addresses	Displays the configured IPv6 addresses.
Default Routers	Displays the default IPv6 Router address(es).

If you change any of the network connectivity parameters, click **Submit** to apply the changes to the system.

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

2.5.7 Network Connection NDP Summary

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

To access this page, click **System > Configuration > Network Connection NDP Summary**.

Figure 2-11: Network NDP Summary

Table 2-11: Network NDP Summary Fields

Field	Description
IPv6 Address	Displays the IP address of the neighbor.
MAC Address	Displays the MAC address of the neighbor.
isRtr	Indicates whether the neighbor is a router. If the neighbor is a router, the value is TRUE. If the neighbor is not a router, the value is FALSE.
Neighbor State	Specifies the state of the neighbor cache entry. Following are the states for dynamic entries in the IPv6 neighbor discovery cache: • Incmp: Address resolution is being performed on the entry. A neighbor solicitation message has been sent to the solicited-node multicast address of the target, but the corresponding neighbor advertisement message has not yet been received. • Reachable: Positive confirmation was received within the last Reachable Time milliseconds that the forward path to the neighbor was functioning properly. While in REACH state, the device takes no special action as packets are sent. • Stale: More than ReachableTime milliseconds have elapsed since the last positive confirmation was received that the forward path was functioning properly. While in STALE state, the device takes no action until a packet is sent. • Delay: More than ReachableTime milliseconds have elapsed since the last positive confirmation was received that the forward path was functioning properly. A packet was sent within the last DELAY_FIRST_PROBE_TIME seconds. If no reachability confirmation is received within DELAY_FIRST_PROBE_TIME seconds of entering the DELAY state, send a neighbor solicitation message and change the state to PROBE. • Probe: A reachability confirmation is actively sought by resending neighbor solicitation messages every RetransTimer milliseconds until a reachability confirmation is received.
Last Updated	Displays the time since the address was confirmed to be reachable.

2.5.8 Network Port DHCPv6 Client Statistics

The Network Port DHCPv6 Client Statistics page displays DHCPv6 client statistics.

To display the Network Port DHCP Client Statistics page, click **System > Configuration > Network Port DHCPv6 Client Statistics**.

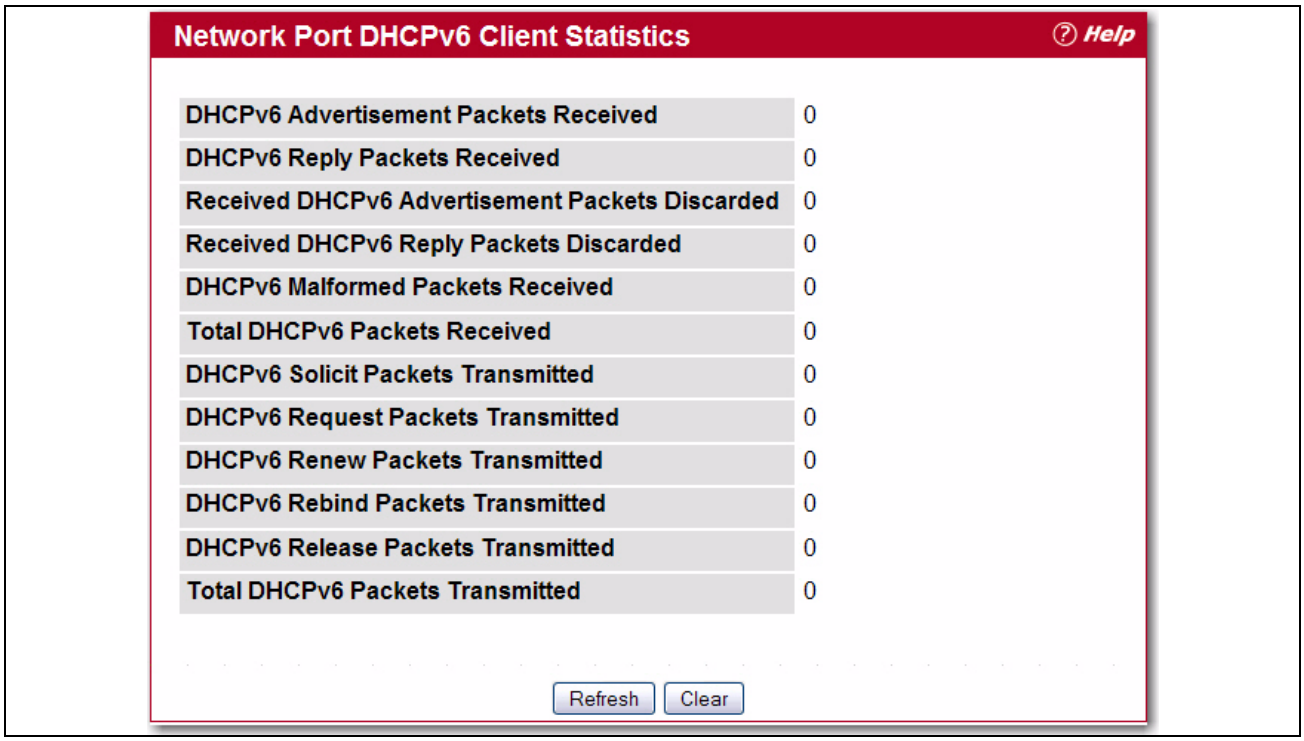

Figure 2-12: Network Port DHCPv6 Client Statistics

Table 2-12: Network Port DHCPv6 Client Statistics Fields

Field	Description
DHCPv6 Advertisement Packets Received	Displays the number of DHCPv6 Advertisement packets received on the network port.
DHCPv6 Reply Packets Received	Displays the number of DHCPv6 Reply packets received on the network port.
Received DHCPv6 Advertisement Packets Discarded	Displays the number of DHCPv6 Advertisement packets discarded on the network port.
Received DHCPv6 Reply Packets Discarded	Displays the number of DHCPv6 Reply packets discarded on the network port.
DHCPv6 Malformed Packets Received	Displays the Number of DHCPv6 packets that are received malformed on the network port.
Total DHCPv6 Packets Received	Displays the total number of DHCPv6 packets received on the network port.
DHCPv6 Solicit Packets Transmitted	Displays the number of DHCPv6 Solicit packets transmitted on the network port.
DHCPv6 Request Packets Transmitted	Displays the number of DHCPv6 Request packets transmitted on the network port.
DHCPv6 Renew Packets Transmitted	Displays the number of DHCPv6 Renew packets transmitted on the network port.
DHCPv6 Rebind Packets Transmitted	Displays the number of DHCPv6 Rebind packets transmitted on the network port.
DHCPv6 Release Packets Transmitted	Displays the number of DHCPv6 Release packets transmitted on the network port.
Total DHCPv6 Packets Transmitted	Displays the total number of DHCPv6 packets transmitted on the network port.

2.5.9 DHCP Client Options

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

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

Figure 2-13: HTTP Configuration

Table 2-13: HTTP Configuration Fields

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

2.5.10 HTTP Configuration

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

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

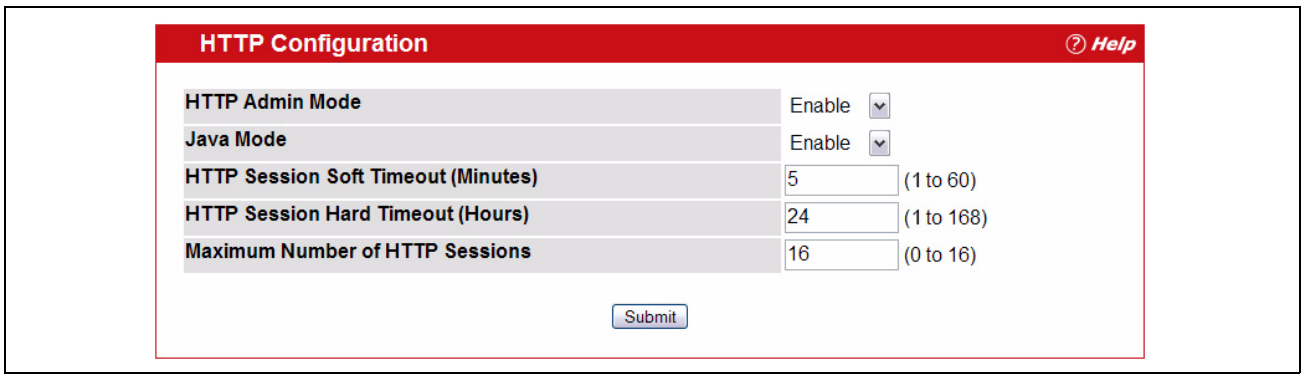

Figure 2-14: HTTP Configuration

Table 2-14: HTTP Configuration Fields

Field	Description
HTTP Admin Mode	This select field is used to Enable or Disable the Administrative Mode of HTTP. The currently configured value is shown when the web page is displayed. The default value is Enable. If you disable the HTTP admin mode, access to the web interface is limited to secure HTTP, which is disabled by default.
Java Mode	This select field is used to Enable or Disable the web Java Mode. This applies to both secure and un-secure HTTP connections. The currently configured value is shown when the web page is displayed. The default value is Enable.
HTTP Session Soft Timeout	This field is used to set the inactivity timeout for HTTP sessions. The value must be in the range of (1 to 60) minutes. A value of zero corresponds to an infinite timeout. The default value is 5 minutes. The currently configured value is shown when the web page is displayed.
HTTP Session Hard Timeout	This field is used to set the hard timeout for HTTP sessions. This timeout is unaffected by the activity level of the session. The value must be in the range of (1 to 168) hours. A value of zero corresponds to an infinite timeout. The default value is 24 hours. The currently configured value is shown when the web page is displayed.
Maximum Number of HTTP Sessions	This field is used to set the maximum allowable number of HTTP sessions. The value must be in the range of (0 to 16). The default value is 16. The currently configured value is shown when the web page is displayed.

If you make changes to the page, click **Submit** to apply the changes to the system.

2.5.11 Telnet Session

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

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

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

To display the Telnet Session Configuration page, click **System > Configuration > Telnet Session** in the navigation tree.

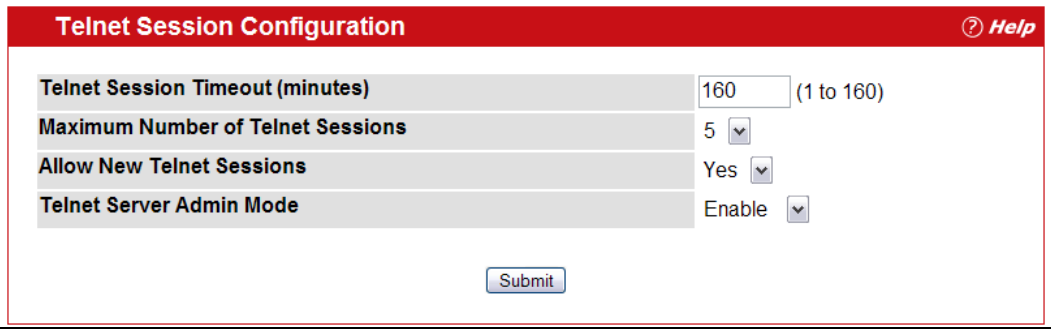

Figure 2-15: Telnet Session Configuration

Table 2-15: Telnet Session Configuration Fields

Field	Description
Telnet Session Timeout (minutes)	Specify how many minutes of inactivity should occur on a telnet session before the session is logged off. You may enter any number from 1 to 160. The factory default is 5. **Note:** When you change the timeout value, the new value is applied to all active and inactive sessions immediately. Any sessions that have been idle longer than the new timeout value are disconnected immediately.
Maximum Number of Telnet Sessions	From the drop-down menu, select how many simultaneous telnet sessions to allow. The maximum is 5, which is also the factory default. A value of 0 indicates that no outbound Telnet session can be established.
Allow New Telnet Sessions	Controls whether to allow new telnet sessions: • **Yes**: Permits new telnet sessions until the maximum number allowed is reached. • **No**: New telnet sessions will not be allowed, but existing sessions are not disconnected.
Telnet Server Admin Mode	Administrative mode for inbound telnet sessions. Setting this value to disable shuts down the telnet port. If the admin mode is set to disable, then all existing telnet connections are disconnected. The default value is Enable.

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

2.5.12 Serial Port

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

To view or configure the serial port settings on the switch, click **System > Configuration > Serial Port Configuration** in the navigation tree.

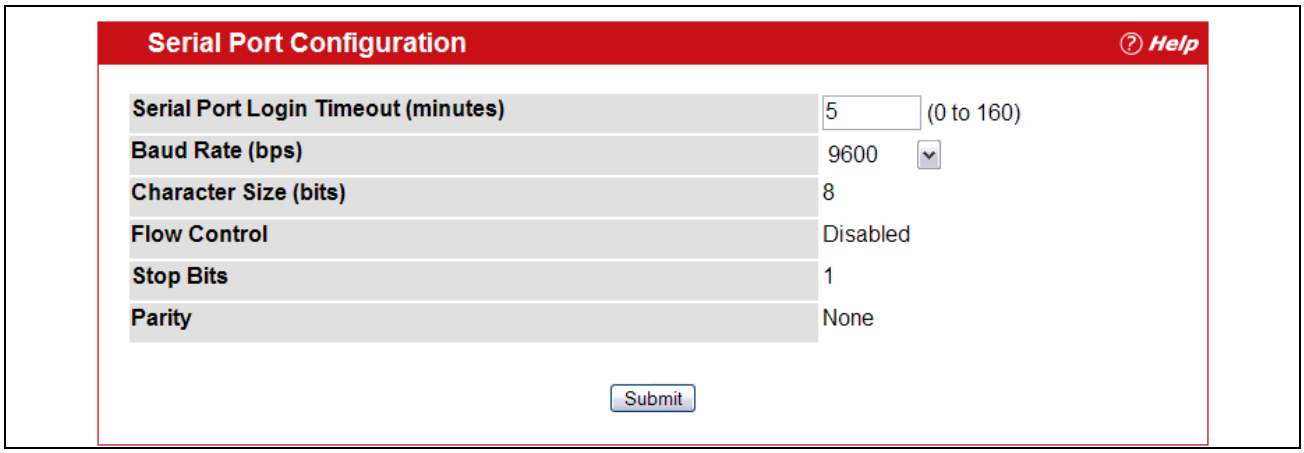

Figure 2-16: Serial Port

Table 2-16: Serial Port Fields

Field	Description
Serial Port Login Timeout (minutes)	Indicates how many minutes of inactivity should occur on a serial port connection before the switch closes the connection. Enter a number between 0 and 160. The factory default is 5. Entering 0 disables the timeout.
Baud Rate (bps)	Select the default baud rate for the serial port connection from the menu. The factory default is 115200 baud for Linux platforms and 9600 baud for VxWorks platforms.
Character Size (bits)	The number of bits in a character. This is always 8.
Flow Control	Whether hardware flow control is enabled or disabled. It is always disabled.
Stop Bits	The number of stop bits per character. Its is always 1.
Parity	The parity method used on the serial port. It is always None.

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

2.5.13 User Accounts

By default, the switch contains two user accounts:

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

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

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

Note...

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

To access the User Accounts page, click **System > Configuration > User Accounts** in the navigation tree.

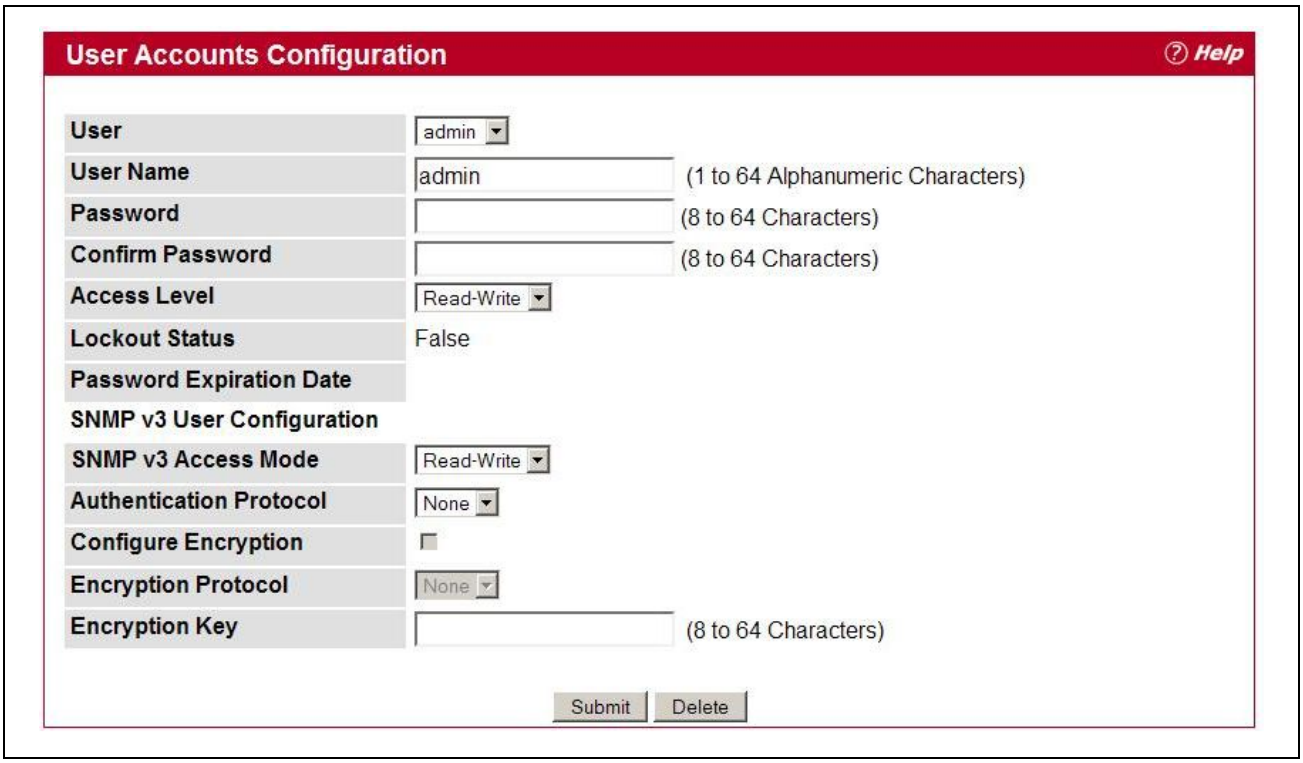

Figure 2-17: User Accounts

Table 2-17: User Accounts Fields

Field	Description
User	From the **User** menu, select an existing user to configure, or select **Create** to create a new user account. The system can have a maximum of five 'Read Only' accounts and one Read/Write account.
User Name	Enter the name you want to give to the new account. (You can only enter data in this field when you are creating a new account.) User names are up to 64 alphanumeric characters in length and are not case sensitive. Valid characters include all the alphanumeric characters and the dash ('-') and underscore ('_') characters. User name *default* is not valid. **Note:** You can change the Read/Write user name from “admin” to something else, but when you click **Submit**, you must re-authenticate with the new username.
Password	Enter the optional new or changed password for the account. It will not display as it is typed, only asterisks (*) or dots(.) will show based on the browser used. Passwords must be greater than eight characters and can be up to 64 characters in length, and are case sensitive.
Confirm Password	Enter the password again, to confirm that you entered it correctly. This field will not display, but will show asterisks (*)
Access Level	Indicates the user's access level. The admin account always has Read/Write access, and all other accounts have Read Only access.
Lockout Status	Indicates whether the user is currently locked out. A user is locked out after a configurable number of failed login attempts. See 2.5.20Password Management 48 for instructions on configuring this number.
Password Expiration Date	Indicates the date when this user’s current password will expire. This is determined by the date the password was created and the number of days specified in the aging Password Aging setting on the Password Management page.
SNMP v3 User Configuration	

Table 2-17: User Accounts Fields (Continued)

Field	Description
SNMP v3 Access Mode	Shows the SNMPv3 access privileges for the user account. The admin account always has 'Read/Write' access, and all other accounts have 'Read Only' access.
Authentication Protocol	Specify the SNMPv3 Authentication Protocol setting for the selected user account. The valid Authentication Protocols are **None**, **MD5** or **SHA**. If you select **None**, the user will be unable to access the SNMP data from an SNMP browser. If you select **MD5** or **SHA**, the user login password will be used as the SNMPv3 authentication password, and you must specify a valid password.
Configure Encryption	Select the check box to change the Encryption Protocol and Encryption Key.
Encryption Protocol	Specify the SNMPv3 Encryption Protocol setting for the selected user account. The valid Encryption Protocols are **None** or **DES**. If you select the **DES** Protocol you must enter a key in the **Encryption Key** field. If **None** is specified for the Protocol, the **Encryption Key** field is not active for input.
Encryption Key	If you selected **DES** in the **Encryption Protocol** field enter the SNMPv3 Encryption Key here. Otherwise this field is not active. Key should be 8 characters in length.

2.5.13.1 Adding a User Account

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

1. From the **User** menu, select **Create**.

   The screen refreshes.

2. Enter a username and password for the new user, then re-enter the password in the **Confirm Password** field.
3. Click **Submit** to update the switch with the values on this screen.

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

2.5.13.2 Changing User Account Information

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

1. From the **User** menu, select the user to change.

   The screen refreshes.

2. To alter the username or, delete the existing name in the **Username** field and enter the new username.

   To change the password, delete any asterisks (*) in the **Password** and **Confirm Password** fields, and then enter and confirm the new password.

3. Click **Submit** to update the switch with the values on this screen.

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

2.5.13.3 Deleting a User Account

Use the following procedures to delete any of the Read Only user accounts.

1. From the **User** menu, select the user to delete.

   The screen refreshes.

2. Click **Delete** to delete the user.

   This button is only visible when you have selected a user account with 'Read Only' access. You cannot delete the 'Read/Write' user.

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

2.5.14 Authentication List Configuration

Use the Authentication List Configuration page to configure login lists. A login list specifies one or more authentication methods to validate switch or port access for the users associated with the list.

Note...

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

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

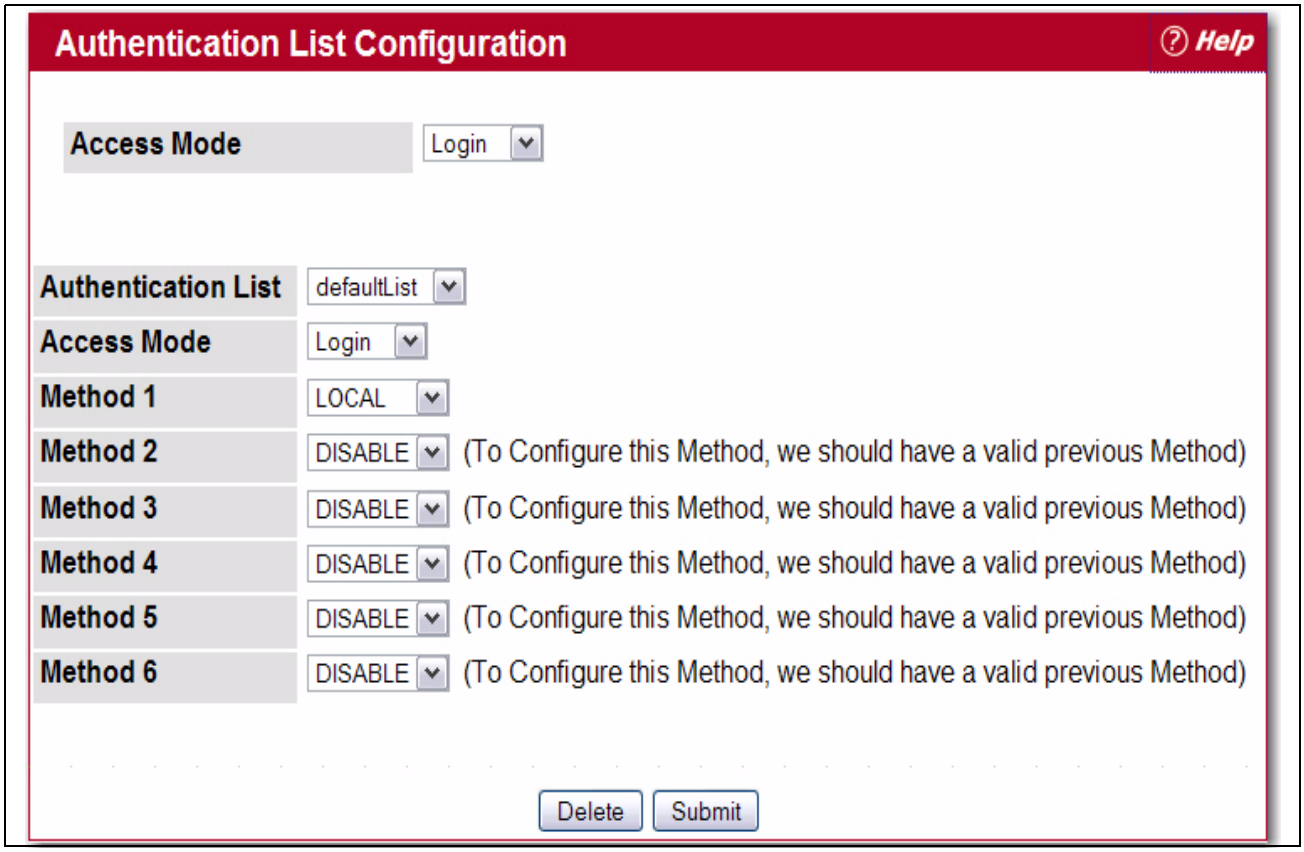

Figure 2-18: Authentication List Configuration

Figure 2-19 shows the fields on the Authentication List Configuration page.

Table 2-18: Authentication List Configuration Fields

Field	Description
Access Mode	A **Login** or **Enable** list specifies the authentication method you want used to validate switch or port access for the users associated with the list. The pre-configured users (admin and guest) are assigned to a pre-configured list named defaultList, which you may not delete. All newly created users are also assigned to the defaultList until you assign them to a different list.
Authentication List	The menu allows you to create a new authentication list or to select an existing list to view or configure.
Method 1	Use the menu to select the method that should appear first in the selected authentication login list. User authentication occurs in the order the methods are selected. Possible methods are as follows: • **enable**: Uses the enable password for authentication. • **line**: Uses the Line password for authentication. • **local**: The user's locally stored ID and password will be used for authentication. • **none**: No authentication is used. **radius:** The user's ID and password will be authenticated using the RADIUS server instead of locally. • **tacacs+**: The user's ID and password will be authenticated using the TACACS+ server. • **undefined**: The authentication method is unspecified. This option cannot be assigned as **Method 1**. ***Note:*** If you select a method that does not time out as the **Method 1** (such as **local**) no other method will be tried, even if you have specified more than one method.
Method 2	Use the menu to select the method, if any, that should appear second in the selected authentication login list. This is the method that will be used if the first method times out. If you select a method that does not time out as the **Method 2**, the **Method 3** will not be tried.
Method 3	Use the menu to select the method, if any, that should appear third in the selected authentication login list. If you select a method that does not time out as the **Method 3**, the **Method 4** will not be tried.
Method 4	Use the menu to select the method, if any, that should appear fourth in the selected authentication login list. If you select a method that does not time out as the **Method 4**, the **Method 5** will not be tried.
Method 5	Use the menu to select the method, if any, that should appear fifth in the selected authentication login list. If you select a method that does not time out as the **Method 5**, the **Method 6** will not be tried.
Method 6	Use the menu to select the method, if any, that should appear sixth in the selected authentication login list.

When **Create** is selected from **Authentication List**, the **Create Authentication List Configuration** page displays.

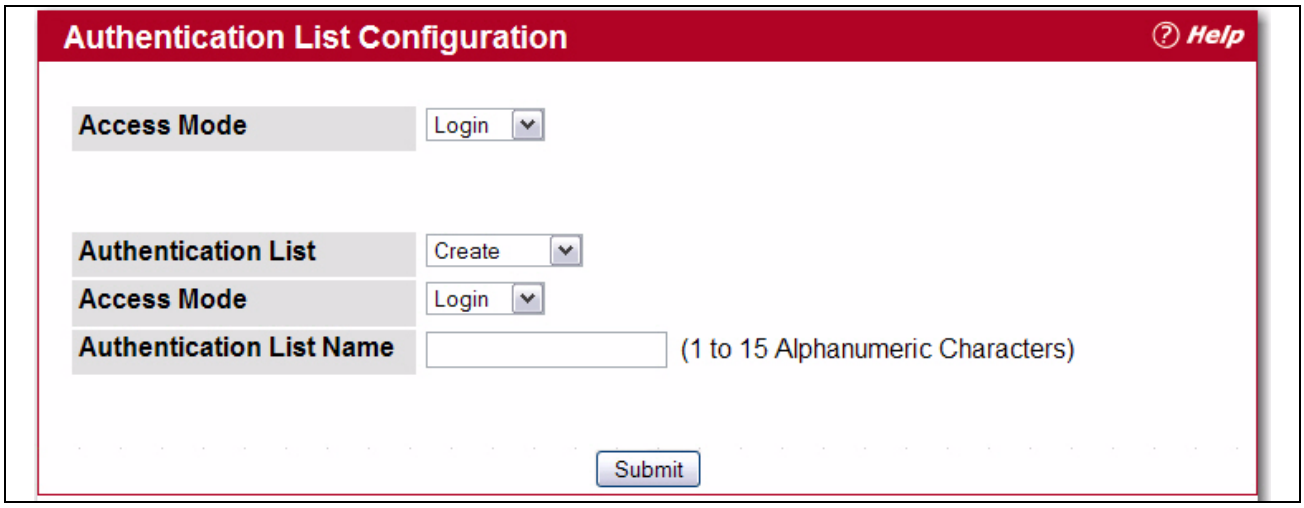

Figure 2-19: Create Authentication List Configuration

Figure 2-19 shows the fields on the Create Authentication List Configuration page.

Table 2-19: Create Authentication List Configuration Fields

Field	Description
Access Mode	A **Login** or **Enable** list specifies the authentication method you want used to validate switch or port access for the users associated with the list. The pre-configured users (admin and guest) are assigned to a pre-configured list named defaultList, which you may not delete. All newly created users are also assigned to the defaultList until you assign them to a different list.
Authentication List	The menu allows you to create a new authentication list or to select an existing list to view or configure.
Authentication List Name	If you are creating a new list, enter the name you want to assign. It can be up to 15 alphanumeric characters long and is not case sensitive.

2.5.14.1 Creating an Authentication List

To create a new authentication list, use the following procedures.

1. Select **Create** from the **Authentication List** field.
2. Select an **Access Mode (Login** or **Enable)** from the drop-down list.
3. In the **Authentication List Name** field, enter a name of 1 to 15 characters.

   The name cannot include spaces.
4. Click **Submit** to create the name and display the Method fields for the new list.

   You are now ready to configure the authentication list. By default, local is set as the initial authentication method.

To retain the changes across a power cycle, you must perform a save.

2.5.14.2 Configuring an Authentication List

To modify an authentication list, use the following procedures.

1. Select an existing list from the **Authentication List** menu.
2. From the **Method 1** field, select the initial login method.

3. If desired, select the second through sixth login method from the **Method** fields.
4. Click **Submit** to apply the changes to the system.

To retain the changes across a power cycle, you must perform a save.

2.5.14.3 Deleting an Authentication List

Use the following procedures to remove an authentication login list from the configuration.

1. Select an existing list from the **Authentication List** menu.
2. Click **Delete**.

   The delete will fail if the selected login list is assigned to any user (including the default user) for system login or IEEE 802.1x port access control. You can only use this button if you have Read/Write access.

To retain the changes across a power cycle, you must perform a save.

2.5.15 Login Session

Use the Login Session page to view information about users who have logged on to the switch.

To access the **Login Session** page, click **System > Configuration > Login Session** in the navigation tree.

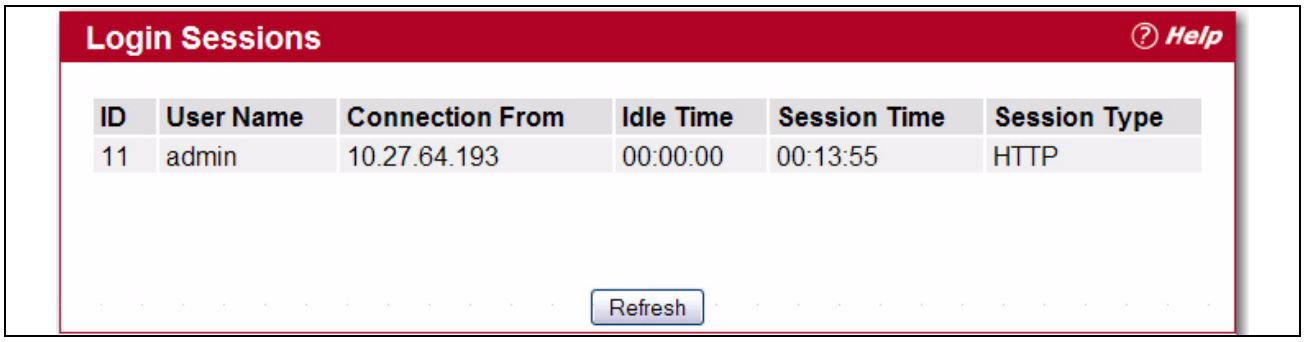

Figure 2-20: Login Session

The Login Session page has the following read-only fields:

Table 2-20: Login Session Fields

Field	Description
ID	Identifies the ID of this row.
User Name	Shows the user name of the user who is currently logged on to the switch.
Connection From	Shows the IP address of the system from which the user is connected. If the connection is a local serial connection, the **Connection From** field entry is EIA-232.
Idle Time	Shows the idle session time.
Session Time	Shows the total session time.
Session Type	Shows the type of session, which can be Telnet, Serial Port, HTTP, or SSH.

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

2.5.16 Authentication List Summary

Use the Authentication List Summary page to view information about the authentication lists on the system.

To access the Authentication List Summary page, click **System > Configuration > Authentication List Summary** in the navigation tree.

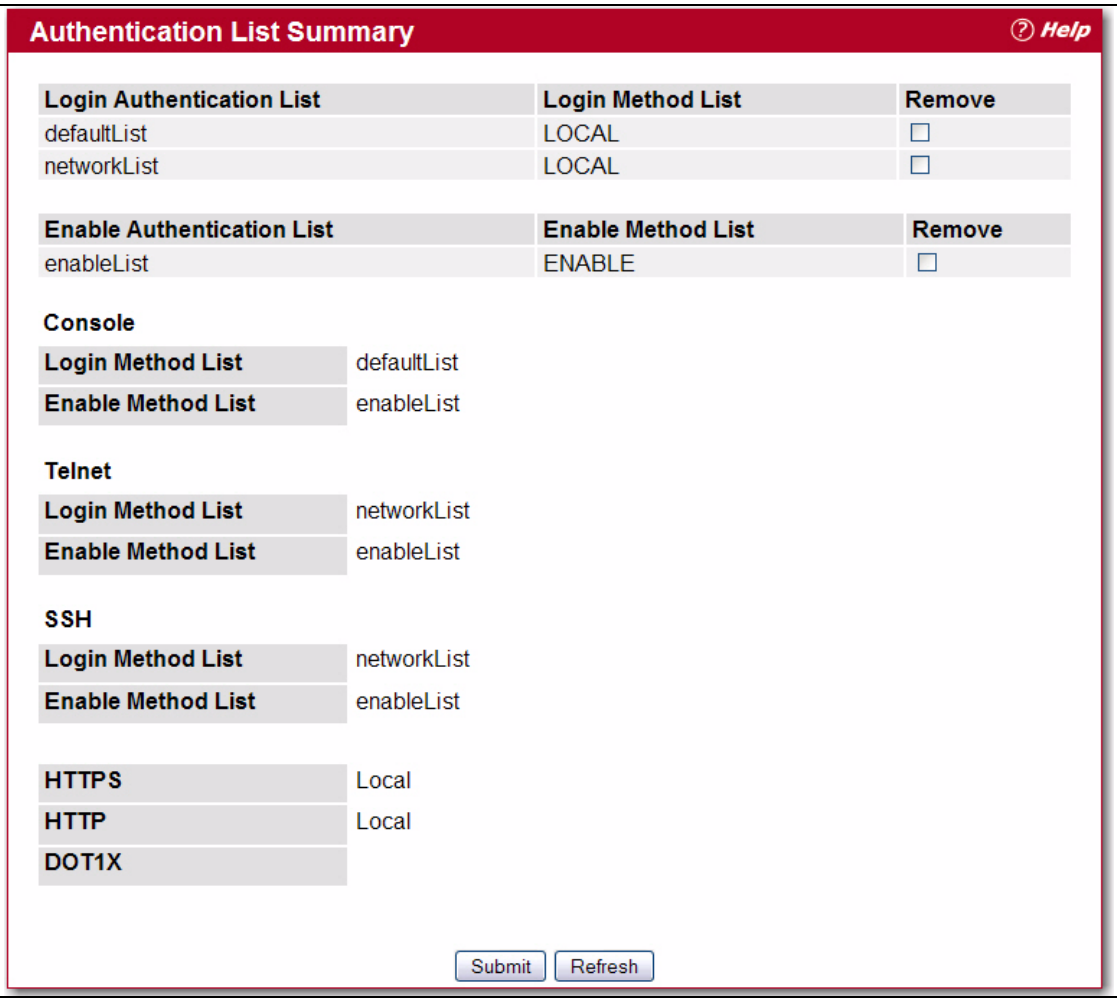

Figure 2-21: Authentication List Summary

Table 2-21 shows the fields for the Authentication List Summary page.

Table 2-21: Authentication List Summary Fields

Field	Description
Login Authentication List	Shows the Login authentication profiles.
Enable Authentication List	Shows the Enable authentication profiles.
Note: Authentication profiles are also shown for: •Console •Telnet •SSH •HTTPS •HTTP •802.1x	
Login/Enable Method List	User authentication methods. Possible options are: • **Enable:** uses the enable password for authentication. • **Line:** uses the Line password for authentication. • **Local:** the user's locally stored ID and password will be used for authentication • **None:** the user is not authenticated **Radius:** the user's ID and password will be authenticated using the RADIUS server instead of locally • **TACACS+:** the user's ID and password will be authenticated using the TACACS+ server.
Remove	Removes the authentication profile when checked.

- Click **Refresh** to update the information on the screen.
- To create a new authentication list, see 2.5.14Authentication List Configuration38. To assign users to a specific authentication list, see 2.5.13User Accounts35. To configure the 802.1x port security users, see RADIUS Settings457.

2.5.17 Select Authentication List

Use the Select Authentication List Configuration page to configure authentication methods for session logins.

To access the Select Authentication List page, click **System > Configuration > Select Authentication List** in the navigation tree.

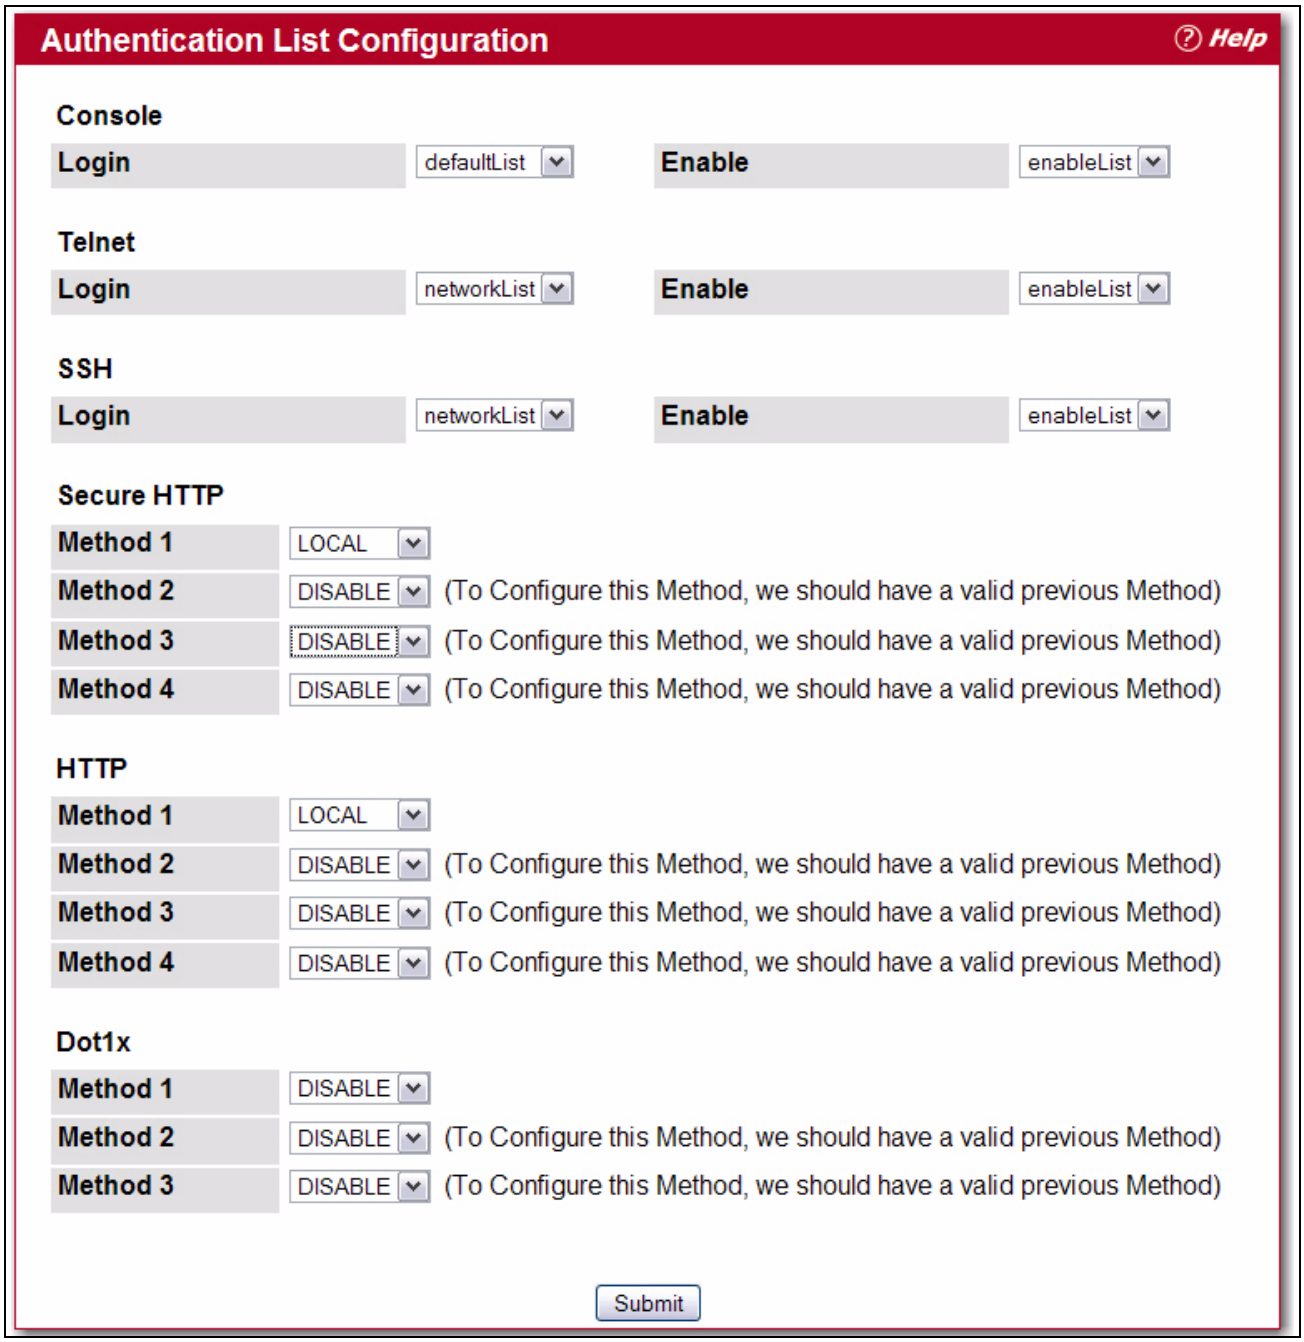

Figure 2-22: Select Authentication List

Table 2-22 shows the fields for the Select Authentication List page.

Table 2-22: Select Authentication List Fields

Field	Description
Console	Authentication profiles used to authenticate console users.

Table 2-22: Select Authentication List Fields (Continued)

Field	Description
Telnet	Authentication profiles used to authenticate Telnet users.
Secure Telnet (SSH)	Authentication profiles used to authenticate Secure Shell (SSH) users. SSH provides clients secure and encrypted remote connections to a device.

HTTP and Secure HTTP	Authentication method used for HTTP access and Secure HTTP access, respectively. Possible field values are:
	• Method 1 - Use the dropdown menu to select the method that should appear first in the selected authentication list. If you select a method that does not time out as the first method, such as 'local' no other method will be tried, even if you have specified more than one method. The options are:
	  • **Enable:** uses the enable password for authentication.
	  • **Line:** uses the Line password for authentication.
	  • **Local:** the user's locally stored ID and password will be used for authentication
	  • **None:** the user is not authenticated
	Radius: the user's ID and password will be authenticated using the RADIUS server instead of locally
	  • **TACACS+:** the user's ID and password will be authenticated using the TACACS+ server
	  • **Reject:** the user is never authenticated
	  • **Undefined:** the authentication method is unspecified (this may not be assigned as the first method)
	• Method 2 - Use the dropdown menu to select the method, if any, that should appear second in the selected authentication list. This is the method that will be used if the first method times out. If you select a method that does not time out as the second method, the third method will not be tried.
	• Method 3 - Use the dropdown menu to select the method, if any, that should appear third in the selected authentication list. This is the method that will be used if the second method times out. If you select a method that does not time out as the third method, the fourth method will not be tried.
	• Method 4 - Use the dropdown menu to select the method, if any, that should appear fourth in the selected authentication list.

Table 2-22: Select Authentication List Fields (Continued)

Field	Description

DOT1X	Authentication method used for Dot1x access. Possible field values are:
	• Method 1 - Use the dropdown menu to select the method that should appear first in the selected authentication list. If you select a method that does not time out as the first method, such as 'local' no other method will be tried, even if you have specified more than one method. The options are:
	  • **Enable:** uses the enable password for authentication.
	  • **Line:** uses the Line password for authentication.
	  • **Local:** the user's locally stored ID and password will be used for authentication
	  • **None:** the user is not authenticated
	Radius: the user's ID and password will be authenticated using the RADIUS server instead of locally
	  • **TACACS+:** the user's ID and password will be authenticated using the TACACS+ server
	  • **Reject:** the user is never authenticated
	  • **Undefined:** the authentication method is unspecified (this may not be assigned as the first method)
	• Method 2 - Use the dropdown menu to select the method, if any, that should appear second in the selected authentication list. This is the method that will be used if the first method times out. If you select a method that does not time out as the second method, the third method will not be tried.
	• Method 3 - Use the dropdown menu to select the method, if any, that should appear third in the selected authentication list.

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

2.5.18 Line Password

Use the Line Password page to configure line mode passwords.

To display the page, click **System > Configuration > Line Password** in the navigation tree.

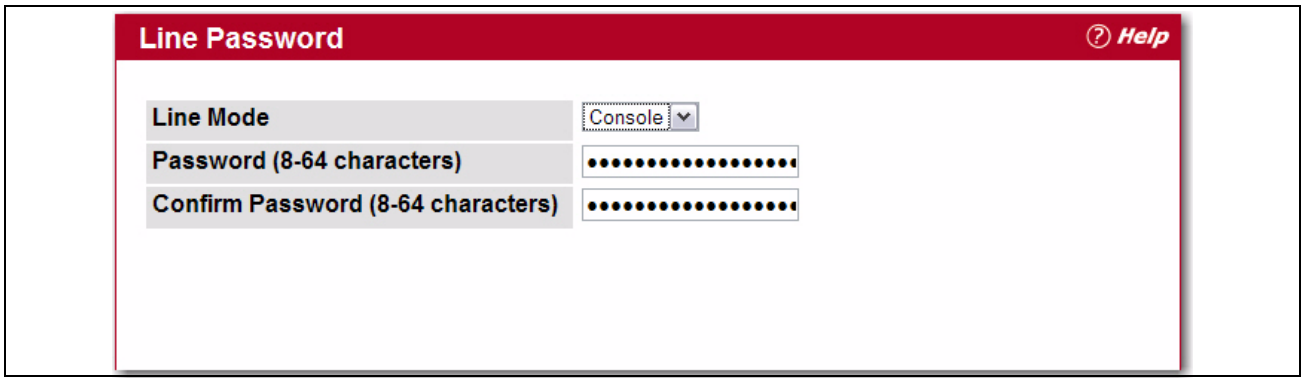

Figure 2-23: Line Password

Table 2-23: Line Password Fields

Field	Description
Line Mode	Select the **Line Mode** from the drop-down list.
Line Password (8-64 characters)	The line password for accessing the device via a console, Telnet, or Secure Telnet session.
Confirm Password (8-64 characters)	Confirms the new line password. The password appears in the ***** format.

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

2.5.19 Enable Password

Use the Enable Password page to configure the enable password.

To display the page, click **System > Configuration > Enable Password** in the navigation tree.

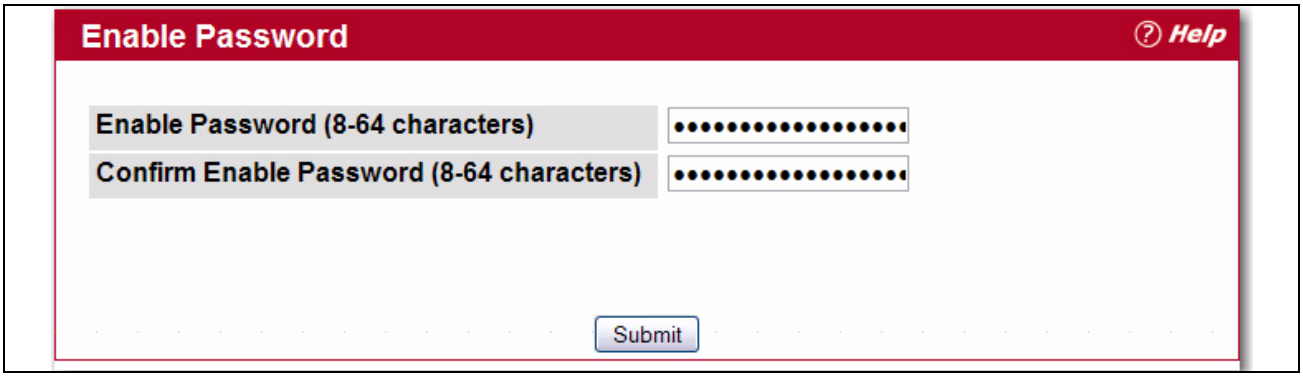

Figure 2-24: Enable Password

Table 2-24: Enable Password Fields

Field	Description
Enable Password (8-64 characters)	The enable password is for accessing the device via a console, Telnet, or Secure Telnet session.
Confirm Enable Password (8-64 characters)	Confirms the new enable password. The password appears in the ***** format.

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

2.5.20 Password Management

Use this page to configure settings that apply to all user passwords.

To display the page, click **Configuration > Password Management** in the navigation tree.

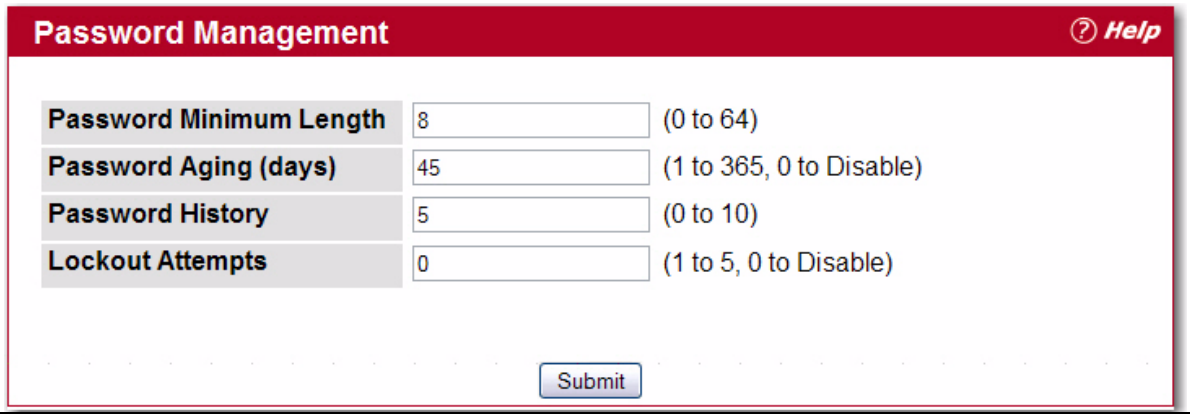

Figure 2-25: Password Management

Table 2-25: Password Management Fields

Field	Description
Password Minimum Length	Passwords must have at least this many characters (8 to 64).
Password Aging (days)	Passwords will expire this many days after creation.
Password History	Users cannot reuse previous passwords up to this number.
Lockout Attempts	After a user fails to log in this number of times, the user is locked out until the password is reset by the administrator.

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

2.5.21 Denial of Service

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

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

Note...

Monitoring and blocking of the types of attacks listed below are supported only on the following platforms:

- BCM56514
- BCM56624
- BCM56820
- BCM56224
- BCM56634
- BCM56636

- **SMAC=DMAC:** Source MAC address=Destination MAC address.
- TCP Port: Source TCP Port = Destination TCP Port.
- UDP Port: Source UDP Port = Destination UDP Port.
- **TCP Flag & Sequence:** TCP Flag SYN set and Source Port < 1024 or TCP Control Flags = 0 and TCP Sequence Number = 0 or TCP Flags FIN, URG, and PSH set and TCP Sequence Number = 0 or TCP Flags SYN and FIN set.
- **TCP Offset:** TCP Header Offset = 1.
- **TCP SYN:** TCP Flag SYN set.
- **TCP SYN & FIN:** TCP Flags SYN and FIN set.
- **TCP FIN & URG & PSH:** TCP Flags FIN and URG and PSH set and TCP Sequence Number = 0.
- **ICMP V6**: Limiting the size of ICMPv6 Ping packets.

- **ICMP Fragment**: Checks for fragmented ICMP packets.

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

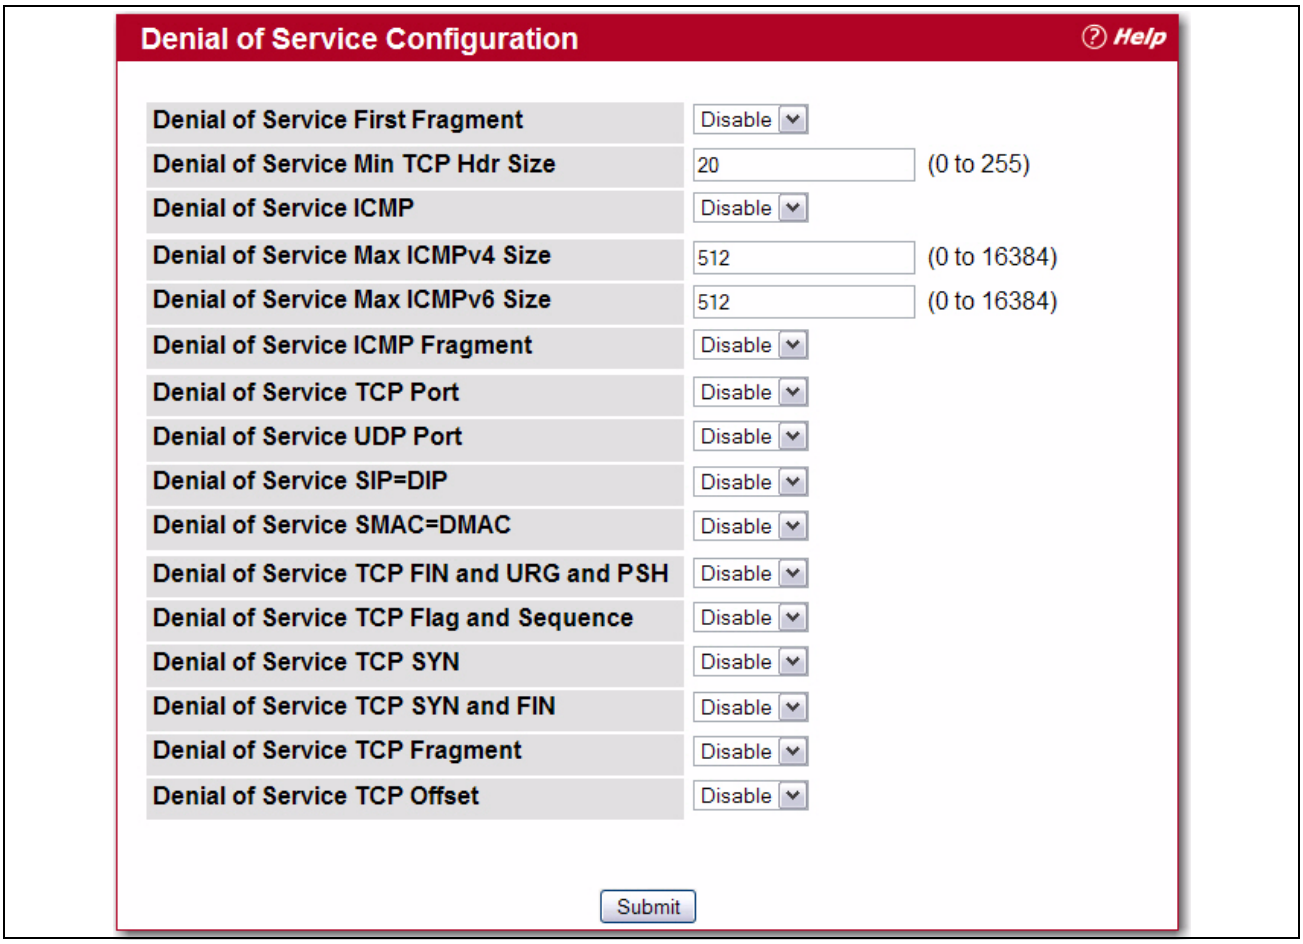

Figure 2-26: Denial of Service

Table 2-26: Denial of Service Configuration Fields

Field	Description
Denial of Service First Fragment	Enable or disable this option by selecting the corresponding line on the pulldown entry field. Enabling First Fragment DoS prevention causes the switch to drop packets that have a TCP header smaller then the configured Min TCP Hdr Size. The factory default is disabled.
Denial of Service Min TCP Hdr Size	Specify the Min TCP Hdr Size allowed. If First Fragment DoS prevention is enabled, the switch will drop packets that have a TCP header smaller then this configured Min TCP Hdr Size. The factory default is disabled.
Denial of Service ICMP	Enable or disable this option by selecting the corresponding line on the pulldown entry field. Enabling ICMP DoS prevention causes the switch to drop ICMP packets that have a type set to ECHO_REQ (ping) and a size greater than the configured ICMP Pkt Size. The factory default is disabled.
Denial of Service Max ICMPv4 Pkt Size	**Note:** This field is only supported on the BCM56514, BCM56624, BCM56820, BCM56224, BCM56634, and BCM56636 platforms. Specify the Max ICMPv4 Pkt Size allowed. If ICMP DoS prevention is enabled, the switch will drop IPv4 ICMP ping packets that have a size greater than this configured Max ICMP Pkt Size. The factory default is disabled.

Table 2-26: Denial of Service Configuration Fields (Continued)

Field	Description
Denial of Service Max ICMPv6 Pkt Size	**Note:** This field is only supported on the BCM56514, BCM56624, BCM56820, BCM56224, BCM56634, and BCM56636 platforms. Specify the Max ICMPv6 ICMP Pkt Size allowed. If ICMP DoS prevention is enabled, the switch will drop IPv6 ICMP ping packets that have a size greater than this configured Max ICMP Pkt Size. The factory default is disabled.
Denial of Service ICMP Fragment	**Note:** This field is only supported on the BCM56514, BCM56624, BCM56820, BCM56224, BCM56634, and BCM56636 platforms. Enable or disable this option by selecting the corresponding line on the pulldown entry field. Enabling ICMP Fragment DoS prevention causes the switch to drop ICMP Fragmented packets. The factory default is disabled.
Denial of Service SIP=DIP	Enable or disable this option by selecting the corresponding line on the pulldown entry field. Enabling SIP=DIP DoS prevention causes the switch to drop packets that have a source IP address equal to the destination IP address. The factory default is disabled.
Denial of Service SMAC=DMAC	**Note:** This field is only supported on the BCM56514, BCM56624, BCM56820, BCM56224, BCM56634, and BCM56636 platforms. Enable or disable this option by selecting the corresponding line on the pulldown entry field. Enabling SMAC=DMAC DoS prevention causes the switch to drop packets that have a source MAC address equal to the destination MAC address. The factory default is disabled.
Denial of Service TCP FIN&URG&PSH	**Note:** This field is only supported on the BCM56514, BCM56624, BCM56820, BCM56224, BCM56634, and BCM56636 platforms. Enable or disable this option by selecting the corresponding line on the pulldown entry field. Enabling TCP FIN & URG & PSH DoS prevention causes the switch to drop packets that have TCP flags FIN, URG, and PSH set and TCP Sequence Number = 0. The factory default is disabled.
Denial of Service TCP Flag&Sequence	**Note:** This field is only supported on the BCM56514, BCM56624, BCM56820, BCM56224, BCM56634, and BCM56636 platforms. Enable or disable this option by selecting the corresponding line on the pulldown entry field. Enabling TCP Flag DoS prevention causes the switch to drop packets that have TCP control flags set to 0 and TCP sequence number set to 0. The factory default is disabled.
Denial of Service TCP Fragment	Enable or disable this option by selecting the corresponding line on the pulldown entry field. Enabling TCP Fragment DoS prevention causes the switch to drop packets that have an IP fragment offset equal to 1. The factory default is disabled.
Denial of Service TCP Offset	**Note:** This field is only supported on the BCM56514, BCM56624, BCM56820, BCM56224, BCM56634, and BCM56636 platforms. Enable or disable this option by selecting the corresponding line on the pulldown entry field. Enabling TCP Offset DoS prevention causes the switch to drop packets that have a TCP header Offset equal to 1. The factory default is disabled.

Table 2-26: Denial of Service Configuration Fields (Continued)

Field	Description
Denial of Service TCP Port	**Note:** This field is only supported on the BCM56514, BCM56624, BCM56820, BCM56224, BCM56634, and BCM56636 platforms. Enable or disable this option by selecting the corresponding line on the pulldown entry field. Enabling TCP Port DoS prevention causes the switch to drop packets that have TCP source port equal to TCP destination port. The factory default is disabled.
Denial of Service TCP SYN	**Note:** This field is only supported on the BCM56514, BCM56624, BCM56820, BCM56224, BCM56634, and BCM56636 platforms. Enable or disable this option by selecting the corresponding line on the pulldown entry field. Enabling TCP SYN DoS prevention causes the switch to drop packets that have TCP Flags SYN set. The factory default is disabled.
Denial of Service TCP SYN&FIN	**Note:** This field is only supported on the BCM56514, BCM56624, BCM56820, BCM56224, BCM56634, and BCM56636 platforms. Enable or disable this option by selecting the corresponding line on the pulldown entry field. Enabling TCP SYN & FIN DoS prevention causes the switch to drop packets that have TCP Flags SYN and FIN set. The factory default is disabled.
Denial of Service UDP Port	**Note:** This field is only supported on the BCM56514, BCM56624, BCM56820, BCM56224, BCM56634, and BCM56636 platforms. Enable or disable this option by selecting the corresponding line on the pulldown entry field. Enabling UDP Port DoS prevention causes the switch to drop packets that have UDP source port equal to UDP destination port. The factory default is disabled.

If you change any of the DoS settings, click **Submit** to apply the changes to the switch. To preserve the changes across a switch reboot, you must perform a save.

2.6 Configuring and Searching the Forwarding Database

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

2.6.1 Configuration

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

To access the Configuration page, click **System > Forwarding Database > Configuration** in the navigation tree.

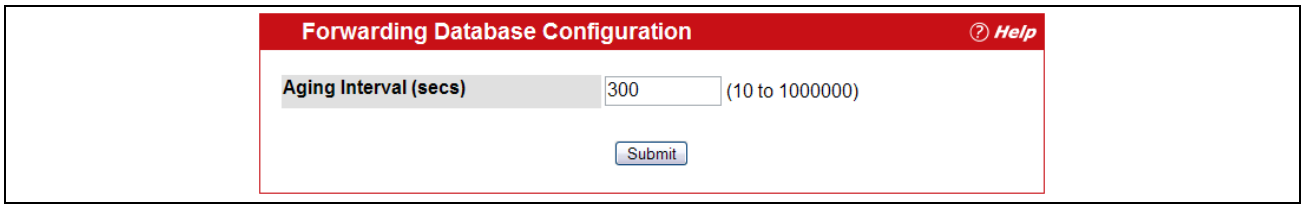

Figure 2-27: Forwarding Database Age-Out Interval

Table 2-27: Forwarding Database Configuration Fields

Field	Description
Aging Interval (secs)	Specify the number of seconds the forwarding database should wait before deleting a learned entry that has not been updated. You may enter any number of seconds between 10 and 1000000.

Note...

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

Click **Submit** to apply the changes to the system. You must perform a save to make the changes persist across a reboot.

2.6.2 Search

Use the Search page to display information about entries in the forwarding database.

To access the Search page, click **System > Forwarding Database > Search** in the navigation tree.

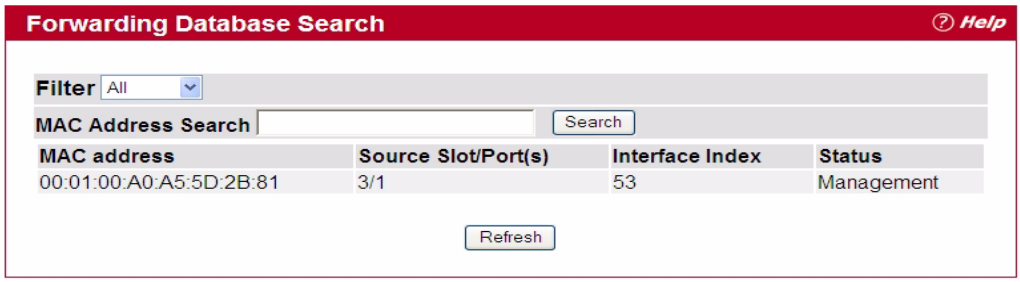

Figure 2-28: Forwarding Database Search

Table 2-28: Forwarding Database Search Fields

Field	Description
Filter	Specify the type of entries to display. When you select a filter from the menu, the screen refreshes and displays the entries based on the filter you select, which can be one of the following: • **Learned**: If you select **Learned**, only MAC addresses that have been learned are displayed. • **All**: If you select **All,** the entire table is displayed.
MAC Address Search	This field allows you to search for an individual MAC address in the forwarding database table.
MAC Address	A unicast MAC address for which the switch has forwarding and/or filtering information. The format is a two byte hexadecimal VLAN ID number followed by a six byte MAC address with each byte separated by colons. For example: 01:23:45:67:89:AB:CD:EF, where 01:23 is the VLAN ID and 45:67:89:AB:CD:EF is the MAC address.
Source Port	The port where this address was learned. In other words, this field shows the port through which the MAC address can be reached.
ifIndex	The ifIndex of the MIB interface table entry associated with the source port.
Status	The status of this entry. The possible values are: • **Static**: The entry was added when a static MAC filter was defined. • **Learned**: The entry was learned by observing the source MAC addresses of incoming traffic, and is currently in use. • **Management**: The system MAC address, which is identified with interface 0.1. • **Self**: The MAC address of one of the switch's physical interfaces.

2.6.2.1 Searching the Forwarding Database

Use the following procedures to search the forwarding database.

1. Enter the two-byte hexadecimal VLAN ID followed by the six byte hexadecimal MAC address in two-digit groups separated by colons.

   For example, 01:23:45:67:89:AB:CD:EF where 01:23 is the VLAN ID and 45:67:89:AB:CD:EF is the MAC address.

2. Click **Search**.

   If the address exists, that entry is displayed as the first entry in the table after the screen refreshes. The entry is followed by the remaining (greater) MAC addresses. An exact match is required. If you click **Refresh**, the MAC addresses with lower values are displayed again.

2.7 Managing Logs

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

The *in-memory* log stores messages in memory based upon the settings for message component and severity.

The Log folder contains links to the following pages:

- Buffered Log Configuration
- Buffered Log
- Command Logger Configuration
- Console Log Configuration
- Event Log
- Hosts Configuration
- Persistent Log Configuration
- Persistent Log
- Syslog Configuration

2.7.1 Buffered Log Configuration

The buffered log stores messages in memory based upon the settings for message component and severity. Use the Buffered Log Configuration page to set the administrative status and behavior of logs in the system buffer.

To access the Buffered Log Configuration page, click **System > Log > Buffered Log Configuration** in the navigation tree.

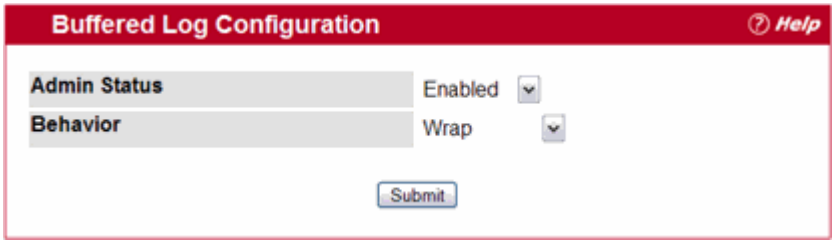

Figure 2-29: Buffered Log Configuration

Table 2-29: Buffered Log Configuration Fields

Field	Description
Admin Status	Determines whether to log messages. • **Enable**: Enables system logging. • **Disable**: Prevents the system from logging messages.
Behavior	Indicates the behavior of the log when it is full. • **Wrap**: When the buffer is full, the oldest log messages are deleted as the system logs new messages. • **Stop on Full**: When the buffer is full, the system stops logging new messages and preserves all existing log messages.

If you change the buffered log settings, click **Submit** to apply the changes to the system. To preserve the changes after a system reboot, you must perform a save.

2.7.2 Buffered Log

Use the Buffered Log page to view the log messages in the system buffer. The newest messages are displayed at the bottom of the page.

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

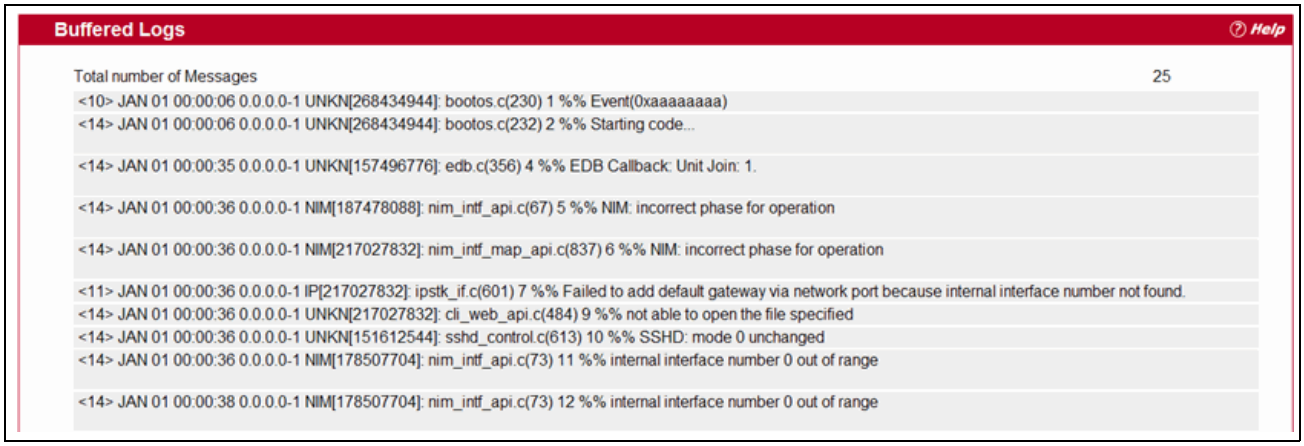

Figure 2-30: Buffered Log

Table 2-30: Buffered Log Fields

Field	Description
Total Number of Messages	Shows the number of buffered messages the system has logged. Only the 128 most recent entries are displayed on the page.

The rest of the page displays the buffered log messages. The following example shows a log message:

```
<15>Aug 24 05:34:05 STK0 MSTP[2110]: mspt_api.c(318) 237 %% Interface 12 transitioned to
root state on message age timer expiry
```

This log message has a severity level of 7 (15 mod 8), which is a debug message.The message was generated by the MSTP component running in thread ID 2110. The message was generated on August 24 05:34:05 by line 318 of file mstp_api.c. This is the 237th message logged.

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

2.7.3 Command Logger Configuration

Use the Command Logger Configuration page to enable the system to log all CLI commands issued on the system. The command log messages are interleaved with the other system logs messages.

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

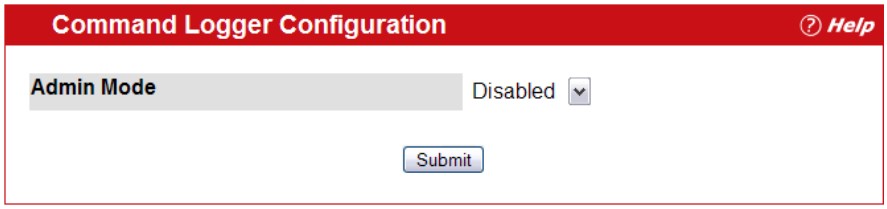

Figure 2-31: Command Logger Configuration

Table 2-31: Command Logger Configuration Fields

Field	Description
Admin Mode	This field determines whether to log CLI commands in the system log file. • **Enable**: The system logs CLI commands. The commands appear in messages on the Buffered Log page. For example, the following log messages shows when the CLI command `show logging buffered` was issued, from which IP address the command was issued, and the name of the user who issued the command: `<5> NOV 29 22:25:00 10.254.24.172-1 UNKN[243420816]: cmd_logger_api.c(87) 34 %% CLI:10.254.24.65:admin:show logging buffered` • **Disable**: This system does not log CLI commands.

If you change the administrative mode, click **Submit** to apply the change to the system.

2.7.4 Console Log Configuration

Use the Console Log Configuration page to control logging to any serial device attached to the switch.

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

Figure 2-32: Console Log Configuration

Table 2-32: Console Log Configuration Fields

Field	Description
Admin Status	From the menu, select whether to enable or disable console logging. The default is disabled. • **Enabled**: Prints log messages to the device attached to the switch serial port. • **Disabled**: Log messages do not print to the device attached to the switch serial port.
Severity Filter	Use the menu to select the severity of the logs to print to the console. Logs with the severity level you select and all logs of greater severity print. For example, if you select Error, the logged messages include Error, Critical, Alert, and Emergency. The default severity level is Alert(1). The severity can be one of the following levels: • **Emergency (0)**: The highest level warning level. If the device is down or not functioning properly, an emergency log is saved to the device. • **Alert (1)**: The second highest warning level. An alert log is saved if there is a serious device malfunction, such as all device features being down. • **Critical (2)**: The third highest warning level. A critical log is saved if a critical device malfunction occurs, for example, two device ports are not functioning, while the rest of the device ports remain functional. • **Error (3)**: A device error has occurred, such as if a port is offline. • **Warning (4)**: The lowest level of a device warning. • **Notice (5)**: Provides the network administrators with device information. • **Informational (6)**: Provides device information. • **Debug (7)**: Provides detailed information about the log. Debugging should only be entered by qualified support personnel.

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

2.7.5 Event Log

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

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

Event Log Help

Entry		Filename	Line	TaskID	Code	Time
00001:	EVENT>	bootos.c	234	0FFFFE00	AAAAAAAA	0 0 0 6
00002:	EVENT>	bootos.c	234	0FFFFE00	AAAAAAAA	0 0 0 6
00003:	EVENT>	bootos.c	234	0FFFFE00	AAAAAAAA	0 0 0 6
00004:	EVENT>	bootos.c	234	0FFFFE00	AAAAAAAA	0 0 0 6
00005:	EVENT>	bootos.c	234	0FFFFE00	AAAAAAAA	0 0 0 6
00006:	ERROR>	unitmgr.c	3641	0FFFFE00	00000000	0 0 0 3
00007:	EVENT>	bootos.c	234	0FFFFE00	AAAAAAAA	0 0 0 6
00008:	EVENT>	unitmgr.c	3641	0CAFFFE8	00000000	0 0 45 3
00009:	EVENT>	bootos.c	234	0FFFFE00	AAAAAAAA	0 0 0 6
00010:	EVENT>	bootos.c	234	0FFFFE00	AAAAAAAA	0 0 0 6
00011:	ERROR>	unitmgr.c	3641	0FFFFE00	00000000	0 0 0 3
00012:	EVENT>	bootos.c	234	0FFFFE00	AAAAAAAA	0 0 0 6
00013:	EVENT>	unitmgr.c	3641	0CB09020	00000000	0 1 42 20
00014:	EVENT>	bootos.c	234	0FFFFE00	AAAAAAAA	0 0 0 6
00015:	EVENT>	unitmgr.c	3641	0CB04330	00000000	0 0 34 58
00016:	EVENT>	bootos.c	234	0FFFFE00	AAAAAAAA	0 0 0 6
00017:	EVENT>	bootos.c	234	0FFFFE00	AAAAAAAA	0 0 0 6
00018:	EVENT>	bootos.c	234	0FFFFE00	AAAAAAAA	0 0 0 6
00019:	EVENT>	bootos.c	234	0FFFFE00	AAAAAAAA	0 0 0 6

Figure 2-33: Event Log

Table 2-33: Event Log Fields

Field	Description
Entry	The number of the entry within the event log. The most recent entry is first.
Filename	The FASTPATH source code filename identifying the code that detected the event.
Line	The line number within the source file of the code that detected the event.
Task ID	The OS-assigned ID of the task reporting the event.
Code	The event code passed to the event log handler by the code reporting the event.
Time	The time the event occurred, measured from the previous reset.

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

2.7.6 Hosts Configuration

Use the Host Configuration page to configure remote logging hosts where the switch can send logs. To enable remote logging, see 2.7.9Syslog Configuration64.

To access the Host Configuration page, click **System > Log > Host Configuration** in the navigation tree.

Figure 2-34 shows the Host Configuration page.

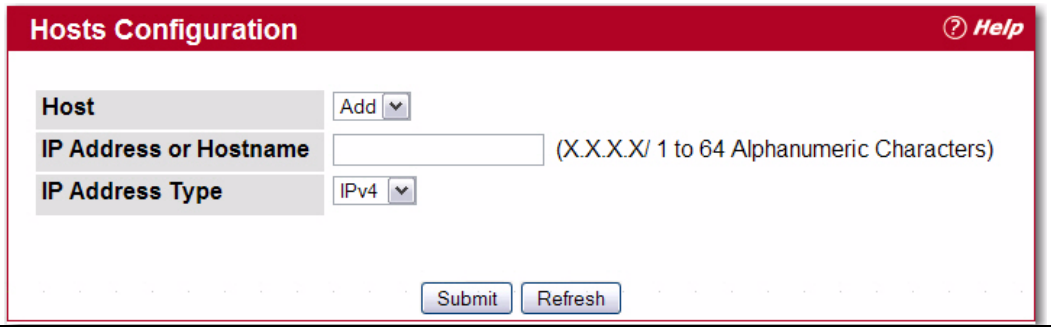

Figure 2-34: Host Configuration

After you add a logging host, the screen displays additional fields, as Figure 2-35 shows.

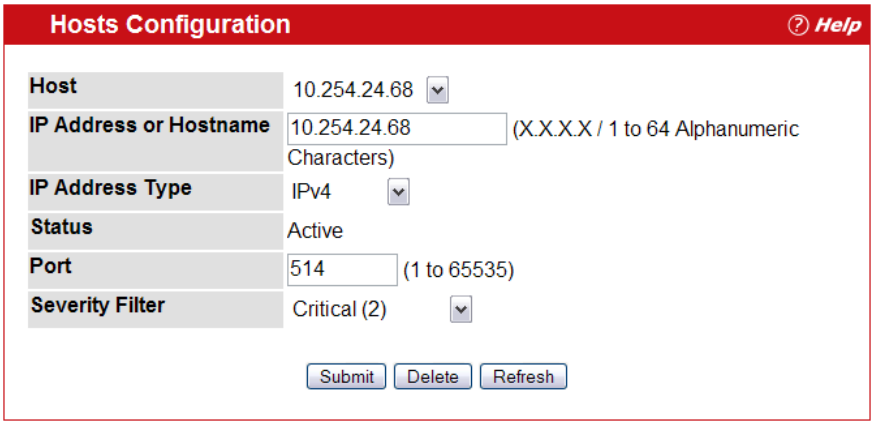

Figure 2-35: Host Configuration with Logging Host

Table 2-34: Host Configuration Fields

Field	Description
Host	Select a host from the list of hosts that have been configured to receive log messages. Select **Add** to add a new host, or select the IP address of an existing host to view or change the settings.
IP Address or Hostname	Enter the IP address or hostname of the host configured for syslog.

Table 2-34: Host Configuration Fields (Continued)

Field	Description
IP Address Type	Select the form of the address entered above: • IPv4: The address is specified in standard dot notation. • DNS: The address is specified as a host name.
Status	Shows whether the remote logging host is currently active.
Port	Identifies the port on the host to which syslog messages are sent. The default port is 514. Specify the port in the text field.
Severity Filter	Use the menu to select the severity of the logs to send to the logging host. Logs with the selected severity level and all logs of greater severity are sent to the host. For example, if you select Error, the logged messages include Error, Critical, Alert, and Emergency. The default severity level is Alert(1). The severity can be one of the following levels: • **Emergency (0)**: The highest level warning level. If the device is down or not functioning properly, an emergency log is saved to the device. • **Alert (1)**: The second highest warning level. An alert log is saved if there is a serious device malfunction, such as all device features being down. • **Critical (2)**: The third highest warning level. A critical log is saved if a critical device malfunction occurs, for example, two device ports are not functioning, while the rest of the device ports remain functional. • **Error (3)**: A device error has occurred, such as if a port is offline. • **Warning (4)**: The lowest level of a device warning. • **Notice (5)**: Provides the network administrators with device information. • **Informational (6)**: Provides device information. • **Debug (7)**: Provides detailed information about the log. Debugging should only be entered by qualified support personnel.

2.7.6.1 Adding a Remote Logging Host

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

1. From the **Host** field, select **Add** to add a new host, or select the IP address of an existing host to configure the host.

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

2.7.6.2 Deleting a Remote Logging Host

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

2.7.7 Persistent Log Configuration

The persistent log is stored in persistent storage, which means that the log messages are retained across a switch reboot.

Note...

Some platforms do not support persistent logging. If your system does not support persistent logging, links to the Persistent Log Configuration and Persistent Log page are not present in the navigation tree menu.

- The first log type is the ***system startup log***. The system startup log stores the first N messages received after system reboot. This log always has the log full operation attribute set to stop on full and can store up to 32 messages.
- The second log type is the ***system operation log***. The system operation log stores the last N messages received during system operation. This log always has the log full operation attribute set to overwrite. This log can store up to 1000 messages.

Either the system startup log or the system operation log stores a message received by the log subsystem that meets the storage criteria, but not both. In other words, on system startup, if the startup log is configured, it stores messages up to its limit. The operation log, if configured, then begins to store the messages.

The system keeps up to three versions of the persistent logs, named <FILE>1.txt, <FILE>2.txt, and <FILE>3.txt. Upon system startup, <FILE>3.txt is removed, <FILE>2.txt is renamed <FILE>3.txt, <FILE>1.txt is renamed <FILE>2.txt, <FILE>1.txt is created and logging begins into <FILE>1.txt. (Replace <FILE> in the above example to specify `olog` for the operation log and `slog` for the startup log.)

The local persistent logs can be retrieved via the Web or CLI, or via xmodem over the local serial cable.

Use the Persistent Log Configuration page to enable or disable persistent logging and to set the severity filter.

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

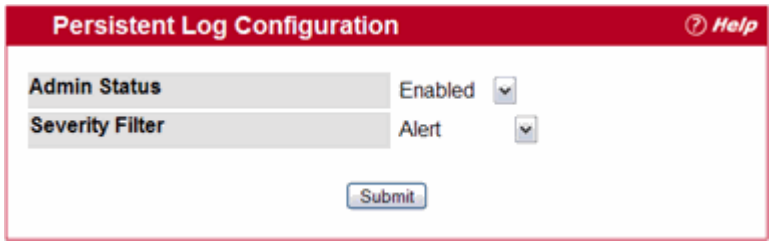

Figure 2-36: Persistent Log Configuration

Table 2-35: Persistent Log Configuration Fields

Field	Description
Admin Status	Select whether to enable or disable persistent logging. The default is disabled. • **Enabled**: Prints log messages to the device attached to the switch serial port. • **Disabled**: Log messages do not print to the device attached to the switch serial port.
Severity Filter	Use the menu to select the severity of the logs to print to the console. Logs with the severity level you select and all logs of greater severity print. For example, if you select Error, the logged messages include Error, Critical, Alert, and Emergency. The default severity level is Alert(1). The severity can be one of the following levels: • **Emergency (0)**: The highest level warning level. If the device is down or not functioning properly, an emergency log is saved to the device. • **Alert (1)**: The second highest warning level. An alert log is saved if there is a serious device malfunction, such as all device features being down. • **Critical (2)**: The third highest warning level. A critical log is saved if a critical device malfunction occurs, for example, two device ports are not functioning, while the rest of the device ports remain functional. • **Error (3)**: A device error has occurred, such as if a port is offline. • **Warning (4)**: The lowest level of a device warning. • **Notice (5)**: Provides the network administrators with device information. • **Informational (6)**: Provides device information. • **Debug (7)**: Provides detailed information about the log. Debugging should only be entered by qualified support personnel.

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

2.7.8 Persistent Log

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

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

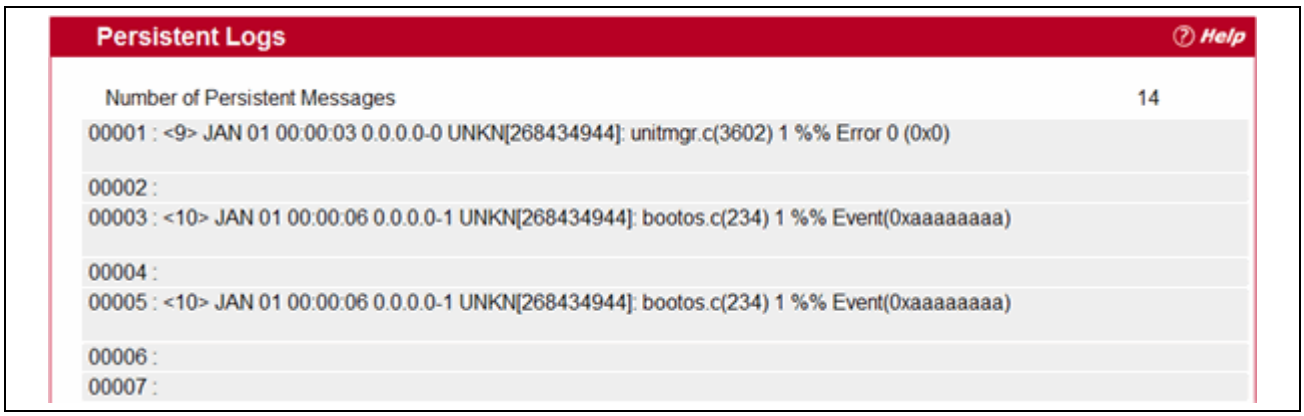

Figure 2-37: Persistent Log

Table 2-36: Persistent Log Fields

Field	Description
Total Number of Messages	Shows the number of persistent messages the system has logged.

The rest of the page displays the log messages. The following example shows a log message:

```
<15>Aug 24 05:34:05 STK0 MSTP[2110]: mspt_api.c(318) 237 %% Interface 12 transitioned to
root state on message age timer expiry
```

This log message has a severity level of 7 (15 mod 8), which is a debug message. The message was generated by the MSTP component running in thread ID 2110. The message was generated on August 24 05:34:05 by line 318 of file mstp_api.c. This is the 237th message logged.

2.7.9 Syslog Configuration

Use the Syslog Configuration page to allow the switch to send log messages to the remote logging hosts configured on the system.

To access the Syslog Configuration page, click **System > Log > Syslog Configuration** in the navigation tree.

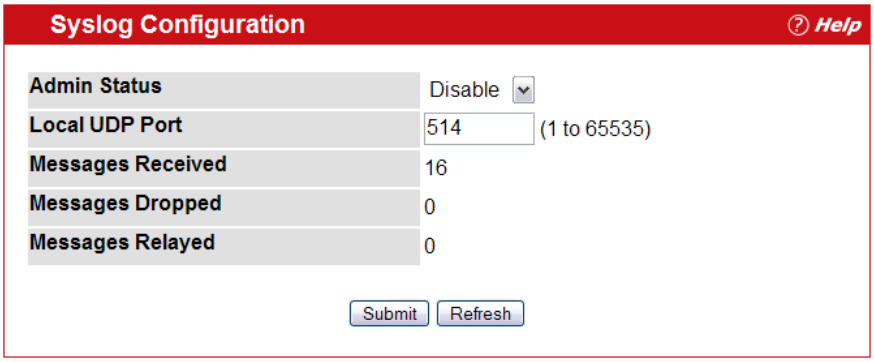

Figure 2-38: Syslog Configuration

Table 2-37: Syslog Configuration Fields

Field	Description
Admin Status	Specifies whether to send log messages to the remote syslog hosts configured on the switch: • **Enable**: Messages will be sent to all configured hosts (syslog collectors or relays) using the values configured for each host. For information about syslog host configuration, see 2.7.6Hosts Configuration59. • **Disable**: Stops logging to all syslog hosts. Disable means no messages will be sent to any collector/relay.

Table 2-37: Syslog Configuration Fields (Continued)

Field	Description
Local UDP Port	Specifies the port on the switch from which syslog messages are sent. The default port is 514.
Messages Received	The number of messages received by the log process. This includes messages that are dropped or ignored.
Messages Dropped	The number of messages that could not be processed due to error or lack of resources.
Messages Relayed	The number of messages forwarded by the syslog function to a syslog host. Messages forwarded to multiple hosts are counted once for each host.

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

2.8 Configuring and Viewing Device Slot Information

The pages in the Slot folder provide information about the cards installed in the slots on the switch. The physical location of the slots depends on the hardware on which FASTPATH software is running. From the Configuration page, you can also manually configure information about cards on some platforms.

2.8.1 Configuration

Use the Card Configuration page to view information about the cards installed in a switch. On some platforms, you can manually configure information about slots.

To access the Card Configuration page, click **System > Slot > Card Configuration** in the navigation menu.

Figure 2-39 shows the fields that display when the slot contains a card.

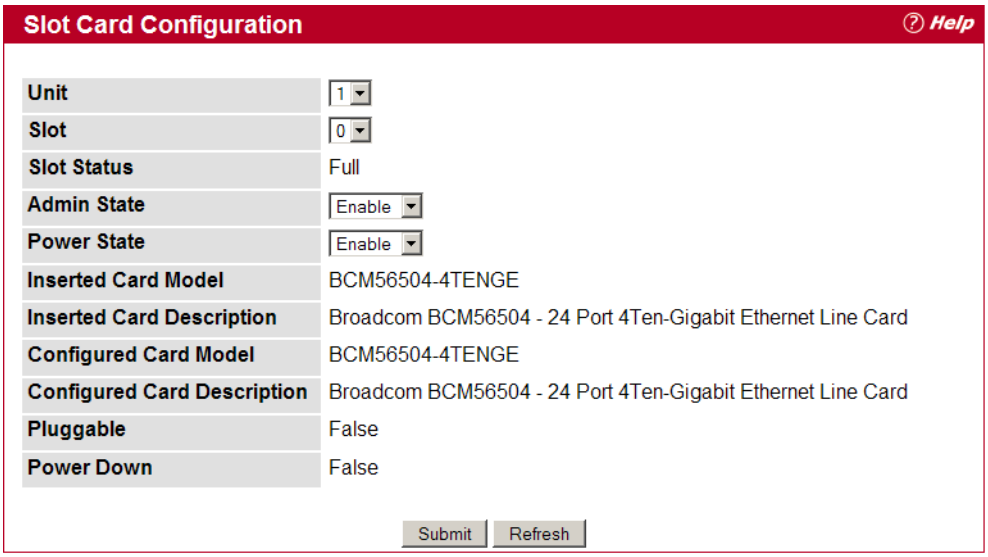

Figure 2-39: Card Configuration

Table 2-38: Card Configuration Fields

Field	Description
Slot	Indicates the slot in the selected slot for which data is to be displayed or configured.
Slot Status	Indicates whether a card is in the slot (**Full** or **Empty**).
Admin State	Displays whether the slot is administratively enabled or disabled. This field is non-configurable for read-only users.
Power State	Displays whether the slot is powered on of off. This field is non-configurable for read-only users.
Card Type	Displays a list of possible supported card types which can be plugged into the slot. This is visible only for slots which do not have any cards plugged into them and which have not already been pre-configured. This field is not visible to read-only users.
Inserted Card Model	Displays the model identifier of the card plugged into the selected slot. If no card has been plugged in, this field is not shown.
Inserted Card Description	Displays the description of the card plugged into the selected slot. If no card has been plugged in, this field is not shown.
Configured Card Model	Displays the model identifier of the card pre-configured for the selected slot. If no card has been pre-configured, this field is not shown.
Configured Card Description	Displays the model identifier of the card pre-configured for the selected slot. If no card has been pre-configured, this field is not shown.
Pluggable	Displays the pluggable indicator of the specified slot.
Power Down	Displays the power down indicator of the specified slot.

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

2.8.2 Slot Summary

The Slot Summary page displays information about the different slots.

To access the Slot Summary page, click **System > Slot > Slot Summary** in the navigation tree.

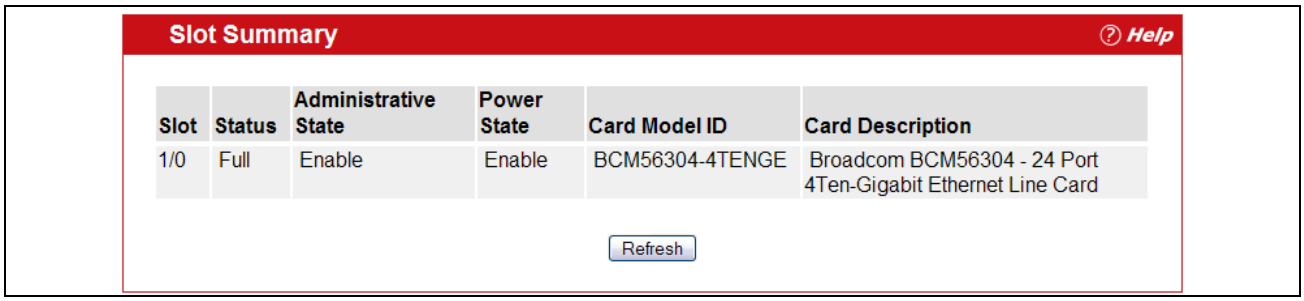

Figure 2-40: Slot Summary

Table 2-39: Slot Summary Fields

Field	Description
Slot	Identifies the slot using the format unit/slot.
Status	Displays whether the slot is empty or full.
Administrative State	Displays whether the slot is administratively enabled or disabled
Power State	Displays whether the slot is powered on of off.
Card Model ID	Displays the model ID of the card configured for the slot.
Card Description	Displays the description of the card configured for the slot.

Click **Refresh** to redisplay the most current information from the router.

2.8.3 Supported Cards

The Supported Cards page provides information about the cards that your platform supports.

To access the Supported Cards page, click **System > Slot > Supported Cards** in the navigation menu.

Supported Cards Help

Card Index	Supported Card	Card Type	Card Model	Card Description
7	BCM56304-4TENGE	0x56304001	BCM56304-4TENGE	Broadcom BCM56304 - 24 Port 4Ten-Gigabit Ethernet Line Card
8	BCM56314-4TENGE	0x56314001	BCM56314-4TENGE	Broadcom BCM56314 - 24 Port 4Ten-Gigabit Ethernet Line Card
9	BCM56504-4TENGE	0x56504001	BCM56504-4TENGE	Broadcom BCM56504 - 24 Port 4Ten-Gigabit Ethernet Line Card
10	BCM56504-4TENGE x2	0x56504401	BCM56504-4TENGE x2	Broadcom BCM56504 - 48 Port 4Ten-Gigabit Ethernet Line Card
11	BCM56514-24GIG-4TENGE	0x56514001	BCM56514-24GIG-4TENGE	Broadcom BCM56514 - 24 Port 4Ten-Gigabit Ethernet Line Card
12	BCM56514-48GIG-4TENGE	0x56514101	BCM56514-48GIG-4TENGE	Broadcom BCM56514 - 48 Port 4Ten-Gigabit Ethernet Line Card
13	BCM56800-20TENGE	0x56800001	BCM56800-20TENGE	Broadcom BCM56800 - 20 Port Ten-Gigabit Ethernet Line Card

Figure 2-41: Supported Cards

Table 2-40: Supported Card Fields

Field	Description
Card Index	Displays the index assigned to the selected card type.
Supported Cards	The menu contains the list of all cards that the system can support. To view information about a card, select it from the drop-down list. The screen refreshes, and the information about that card appears in the other fields on the page.
Card Type	Displays the hardware type of this supported card. This is a 32-bit data field.
Card Model ID	Displays the string to identify the model of the supported card.
Card Descriptor	Displays a data field used to identify the supported card.

Click **Refresh** to redisplay the most current information from the router.

2.9 Configuring and Viewing Device Port Information

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

- Configuration
- Summary
- Port Description
- Cable Test
- Multiple Port Mirroring

2.9.1 Configuration

Use the Port Configuration page to configure the physical interfaces on the switch.

To access the Port Configuration page, click **System > Port > Configuration** in the navigation tree.

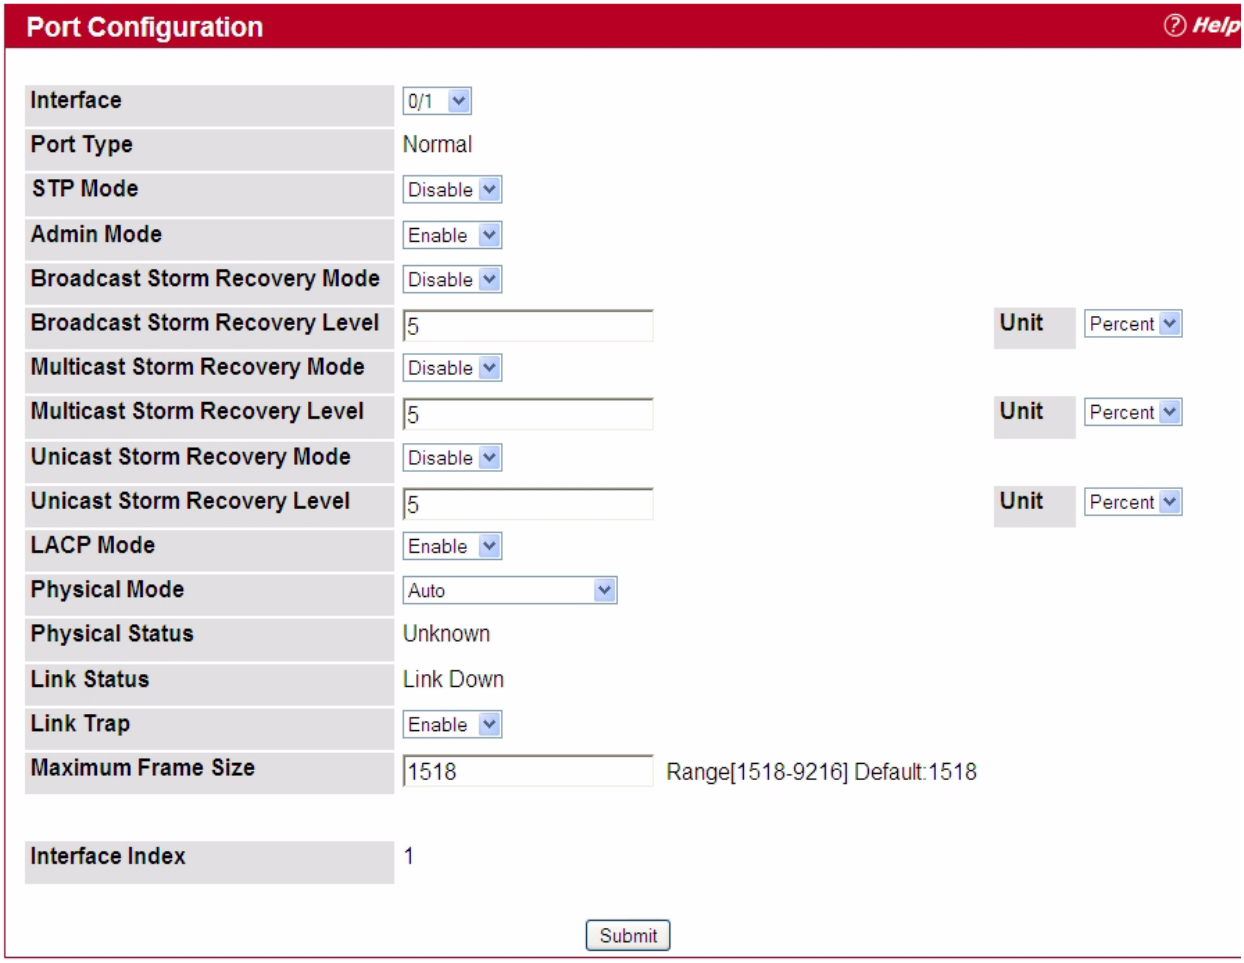

Figure 2-42: Port Configuration

Table 2-41: Port Configuration Fields

Field	Description
Slot/Port	Select the port from the menu to display or configure data for that port. If you select **All**, the changes you make to the **Port Configuration** page apply to all physical ports on the system.
Port Type	For most ports this field is blank. Otherwise the possible values are: • **Mirrored**: Indicates that the port has been configured as a monitoring port and is the source port in a port mirroring session. For more information about port monitoring and probe ports, see 2.9.5Multiple Port Mirroring78. • **Probe**: Indicates that the port has been configured as a monitoring port and is the destination port in a port mirroring session. For more information about port monitoring and probe ports, see 2.9.5Multiple Port Mirroring78. • **Port Channel**: Indicates that the port has been configured as a member of a port-channel, which is also known as a link Aggregation Group (LAG). For information about configuring port channels, see Creating Port Channels248.
STP Mode	Shows the Spanning Tree Protocol (STP) Administrative Mode for the port or LAG. For more information about STP, see Configuring Spanning Tree Protocol 257. The possible values for this field are: • **Enable**: Enables the Spanning Tree Protocol for this port. • **Disable**: Disables the Spanning Tree Protocol for this port.
Admin Mode	Use the pulldown menu to select the port control administration state, which can be one of the following: • **Enable**: The port can participate in the network (default). • **Disable**: The port is administratively down and does not participate in the network.
Broadcast Storm Recovery Mode	Enable or disable this option by selecting one of the following options on the pulldown entry field: • **Enable**: When the broadcast traffic on the specified Ethernet port exceeds the configured threshold, the switch blocks (discards) the broadcast traffic. • **Disable**: The port does not block broadcast traffic if traffic on the port exceeds the configured threshold. The factory default is disabled.
Broadcast Storm Recovery Level	Specify the data rate at which storm control activates. The factory default is 5 percent of port speed. The level units can be set to percent or packets-per-second.
Multicast Storm Recovery Mode	Enable or disable this option by selecting one of the following options on the pulldown entry field: • **Enable**: When the multicast traffic on the specified Ethernet port exceeds the configured threshold, the switch blocks (discards) the multicast traffic. • **Disable**: The port does not block multicast traffic if traffic on the port exceeds the configured threshold. The factory default is disabled.
Multicast Storm Recovery Level	Specify the data rate at which storm control activates. The factory default is 5 percent of port speed. The level units can be set to percent or packets-per-second.
Unicast Storm Recovery Mode	Enable or disable this option by selecting one of the following options on the pulldown entry field: • **Enable**: When the unicast traffic on the specified Ethernet port exceeds the configured threshold, the switch blocks (discards) the unicast traffic. • **Disable**: The port does not block unicast traffic if the unicast traffic on the port exceeds the configured threshold. The factory default is disabled.
Unicast Storm Recovery Level	Specify the data rate at which storm control activates. The factory default is 5 percent of port speed. The level units can be set to percent or packets-per-second.
LACP Mode	Selects the Link Aggregation Control Protocol administration state: • **Enable**: Specifies that the port is allowed to participate in a port channel (LAG), which is the default mode. • **Disable**: Specifies that the port cannot participate in a port channel (LAG).

Table 2-41: Port Configuration Fields (Continued)

Field	Description
Physical Mode	Use the pulldown menu to select the port's speed and duplex mode. If the **Slot/Port** field is set to **All** and you apply a physical mode other than **Auto**, the mode is applied to all applicable interfaces only: • **Auto**: The duplex mode and speed will be set by the auto-negotiation process. The port's maximum capability (full duplex and 100 Mbps) will be advertised. • **<Speed> Half Duplex**: The port speeds available from the menu depend on the platform on which the FASTPATH software is running and which port you select. In half-duplex mode, the transmissions are one-way. In other words, the port does not send and receive traffic at the same time. • **<Speed> Full Duplex**: The port speeds available from the menu depend on the platform on which the FASTPATH software is running and which port you select. In half-duplex mode, the transmissions are two-way. In other words, the port can send and receive traffic at the same time.
Physical Status	Indicates the port speed and duplex mode.
Link Status	Indicates whether the Link is up or down.
Link Trap	This object determines whether or not to send a trap when link status changes. The factory default is enabled: • **Enable**: Specifies that the system sends a trap when the link status changes. • **Disable**: Specifies that the system does not send a trap when the link status changes.
Maximum Frame Size	Indicates the maximum Ethernet frame size the interface supports or is configured to support. The frame size includes the Ethernet header, CRC, and payload. (1518 to 9216). The default maximum frame size is 1518.
ifIndex	The ifIndex of the interface table entry associated with this port. If the **Slot/Port** field is set to **All**, this field is blank.

If you make any changes to the page, click **Submit** to apply the changes to the system.

2.9.2 Summary

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

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

The table on the Port Summary page does not fit on one screen. Use the scroll bar at the bottom of the browser to view all the columns on the page. Figure 2-49 shows the first six rows of all the columns on the page. Although the table is split into three separate images in the figure, the columns are continue horizontally across the page.

Port Summary

MST ID 0

MST : CST

Interface	Port Type	STP Mode	Forwarding State	Port Role	Admin Mode
0/1	Normal	Disable	Disabled	Disabled	Enable
0/2	Normal	Disable	Disabled	Disabled	Enable
0/3	Normal	Disable	Disabled	Disabled	Enable
0/4	Normal	Disable	Disabled	Disabled	Enable
0/5	Normal	Disable	Disabled	Disabled	Enable
0/6	Normal	Disable	Disabled	Disabled	Enable
0/7	Normal	Disable	Disabled	Disabled	Enable
0/8	Normal	Disable	Disabled	Disabled	Enable
0/9	Normal	Disable	Disabled	Disabled	Enable
0/10	Normal	Disable	Disabled	Disabled	Enable
0/11	Normal	Disable	Disabled	Disabled	Enable
0/12	Normal	Disable	Disabled	Disabled	Enable
0/13	Normal	Disable	Disabled	Disabled	Enable
0/14	Normal	Disable	Disabled	Disabled	Enable
0/15	Normal	Disable	Disabled	Disabled	Enable
0/16	Normal	Disable	Disabled	Disabled	Enable
1/1	Normal	Disable	Disabled	Disabled	Enable
1/2	Normal	Disable	Disabled	Disabled	Enable
1/3	Normal	Disable	Disabled	Disabled	Enable
1/4	Normal	Disable	Disabled	Disabled	Enable
1/5	Normal	Disable	Disabled	Disabled	Enable
1/6	Normal	Disable	Disabled	Disabled	Enable
1/7	Normal	Disable	Disabled	Disabled	Enable
1/8	Normal	Disable	Disabled	Disabled	Enable
1/9	Normal	Disable	Disabled	Disabled	Enable
1/10	Normal	Disable	Disabled	Disabled	Enable
1/11	Normal	Disable	Disabled	Disabled	Enable
1/12	Normal	Disable	Disabled	Disabled	Enable
1/13	Normal	Disable	Disabled	Disabled	Enable
1/14	Normal	Disable	Disabled	Disabled	Enable

1/15	Normal	Disable	Disabled	Disabled	Enable
1/16	Normal	Disable	Disabled	Disabled	Enable
1/17	Normal	Disable	Disabled	Disabled	Enable
1/18	Normal	Disable	Disabled	Disabled	Enable
1/19	Normal	Disable	Disabled	Disabled	Enable
1/20	Normal	Disable	Disabled	Disabled	Enable
1/21	Normal	Disable	Disabled	Disabled	Enable
1/22	Normal	Disable	Disabled	Disabled	Enable
1/23	Normal	Disable	Disabled	Disabled	Enable
1/24	Normal	Disable	Disabled	Disabled	Enable
1/25	Normal	Disable	Disabled	Disabled	Enable
1/26	Normal	Disable	Disabled	Disabled	Enable
1/27	Normal	Disable	Disabled	Disabled	Enable
1/28	Normal	Disable	Disabled	Disabled	Enable
1/29	Normal	Disable	Disabled	Disabled	Enable
1/30	Normal	Disable	Disabled	Disabled	Enable
1/31	Normal	Disable	Disabled	Disabled	Enable
1/32	Normal	Disable	Disabled	Disabled	Enable

Figure 2-43: Port Summary

Table 2-42: Port Summary Fields

Field	Description
MST ID	If Spanning Tree Protocol is enabled on the switch, you can select the Multiple Spanning Tree instance ID from the list of all currently configured MST ID's to determine the values displayed for the Spanning Tree parameters. Changing the selected MST ID will generate a screen refresh. If STP is disabled, which is the default, the MST ID field shows the static value “CST” instead of a menu.
Slot/Port	Identifies the port that the information in the rest of the row is associated with.
Port Type	For most ports this field is blank. Otherwise, the possible values are: • **Mirrored**: Indicates that the port has been configured as a monitoring port and is the source port in a port mirroring session. For more information about port monitoring and probe ports, see 2.9.5Multiple Port Mirroring78. • **Probe**: Indicates that the port has been configured as a monitoring port and is the destination port in a port mirroring session. For more information about port monitoring and probe ports, see 2.9.5Multiple Port Mirroring78. • **Port Channel**: Indicates that the port has been configured as a member of a port-channel, which is also known as a link Aggregation Group (LAG). For information about configuring port channels, see Creating Port Channels248.
STP Mode	Shows the Spanning Tree Protocol (STP) Administrative Mode for the port or LAG, which can be **Enabled** or **Disabled**. For more information about STP, see Configuring Spanning Tree Protocol257.

Table 2-42: Port Summary Fields (Continued)

Field	Description
Forwarding State	The port's current state Spanning Tree state. This state controls what action a port takes on receipt of a frame. If the bridge detects a malfunctioning port it will place that port into the broken state. The other five states are defined in IEEE 802.1D: • Disabled • Blocking • Listening • Learning • Forwarding • Broken
Port Role	Each MST Bridge Port that is enabled is assigned a Port Role for each spanning tree. The port role will be one of the following values: Root Port, Designated Port, Alternate Port, Backup Port, Master Port, or Disabled Port.
Admin Mode	Shows the port control administration state, which can be one of the following: • **Enabled**: The port can participate in the network (default). • **Disabled**: The port is administratively down and does not participate in the network.
Bcast Storm Mode	Shows whether the Broadcast Storm Recovery Mode, which can be one of the following: • **Enabled**: When the broadcast traffic on the specified Ethernet port exceeds the configured threshold, the switch blocks (discards) the broadcast traffic. • **Disabled**: The port does not block broadcast traffic if traffic on the port exceeds the configured threshold. The factory default is disabled.
Bcast Storm Level	Shows the Broadcast Storm Recovery Level, which is the data rate at which storm control activates. The value is a percentage of port speed and ranges from 0-100. The factory default is 5 percent of port speed.
Mcast Storm Mode	Shows the Multicast Storm Recovery Mode, which is one of the following: • **Enabled**: When the multicast traffic on the specified Ethernet port exceeds the configured threshold, the switch blocks (discards) the multicast traffic. • **Disabled**: The port does not block multicast traffic if traffic on the port exceeds the configured threshold. The factory default is disabled.
Mcast Storm Level	Shows the Multicast Storm Recovery Level, which is the data rate at which storm control activates. The value is a percentage of port speed and ranges from 0-100. The factory default is 5 percent of port speed.
Ucast Storm Mode	Shows the Unicast Storm Recovery Mode, which can be one of the following: • **Enabled**: When the unicast traffic on the specified Ethernet port exceeds the configured threshold, the switch blocks (discards) the unicast traffic. • **Disabled**: The port does not block unicast traffic if the unicast traffic on the port exceeds the configured threshold. The factory default is disabled.
Ucast Storm Level	**Shows the Unicast Storm Recovery Level**, which is the data rate at which storm control activates. The value is a percentage of port speed and ranges from 0-100.The factory default is 5 percent of port speed.
LACP Mode	Indicates the Link Aggregation Control Protocol administration state. The mode must be enabled in order for the port to participate in Link Aggregation. This field can have the following values: • **Enable**: Specifies that the port is allowed to participate in a port channel (LAG), which is the default mode. • **Disable**: Specifies that the port cannot participate in a port channel (LAG).

Table 2-42: Port Summary Fields (Continued)

Field	Description
Physical Mode	Shows the speed and duplex mode at which the port is configured: • **Auto**: The duplex mode and speed will be set by the auto-negotiation process. The port's maximum capability (full duplex and 100 Mbps) will be advertised. • **<Speed> Half Duplex**: The port speeds available from the menu depend on the platform on which the FASTPATH software is running and which port you select. In half-duplex mode, the transmissions are one-way. In other words, the port does not send and receive traffic at the same time. • **<Speed> Full Duplex**: The port speeds available from the menu depend on the platform on which the FASTPATH software is running and which port you select. In half-duplex mode, the transmissions are two-way. In other words, the port can send and receive traffic at the same time.
Physical Status	Indicates the port speed and duplex mode at which the port is operating.
Link Status	Indicates whether the Link is up or down.
Link Trap	This object determines whether or not to send a trap when link status changes. The factory default is enabled. • **Enable**: Specifies that the system sends a trap when the link status changes. • **Disable**: Specifies that the system does not send a trap when the link status changes.

Click **Refresh** to redisplay the most current information from the router.

2.9.3 Port Description

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

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

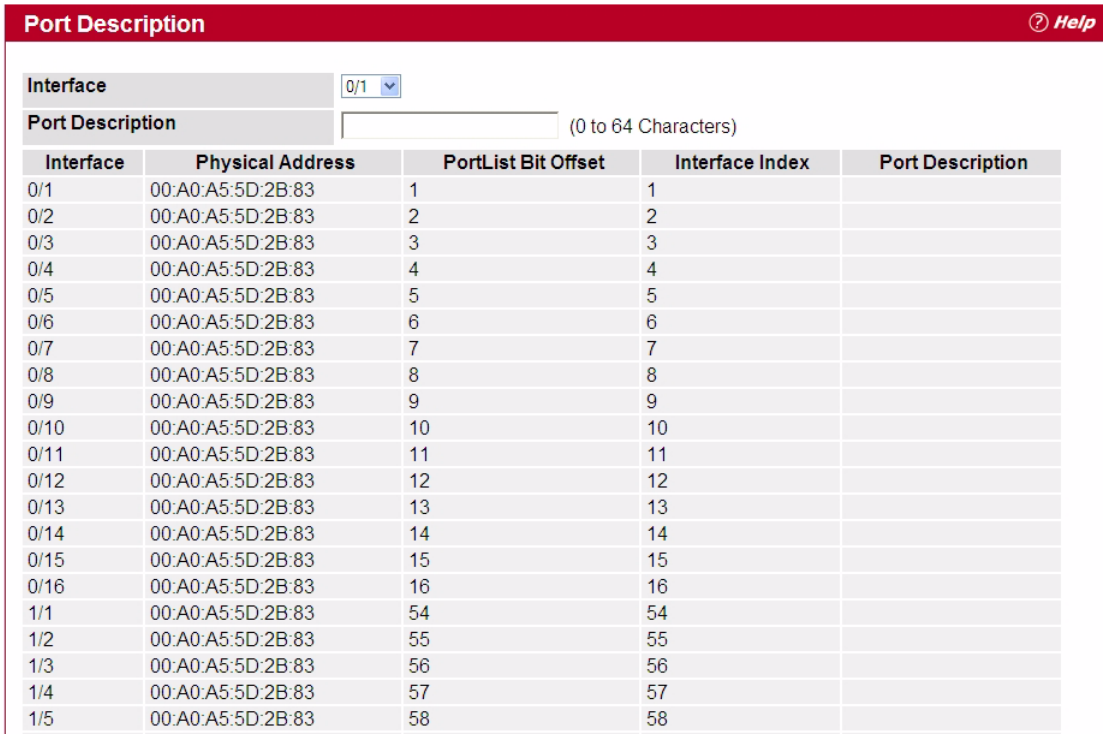

Interface	Physical Address	PortList Bit Offset	Interface Index	Port Description
0/1	00:A0:A5:5D:2B:83	1	1	
0/2	00:A0:A5:5D:2B:83	2	2	
0/3	00:A0:A5:5D:2B:83	3	3	
0/4	00:A0:A5:5D:2B:83	4	4	
0/5	00:A0:A5:5D:2B:83	5	5	
0/6	00:A0:A5:5D:2B:83	6	6	
0/7	00:A0:A5:5D:2B:83	7	7	
0/8	00:A0:A5:5D:2B:83	8	8	
0/9	00:A0:A5:5D:2B:83	9	9	
0/10	00:A0:A5:5D:2B:83	10	10	
0/11	00:A0:A5:5D:2B:83	11	11	
0/12	00:A0:A5:5D:2B:83	12	12	
0/13	00:A0:A5:5D:2B:83	13	13	
0/14	00:A0:A5:5D:2B:83	14	14	
0/15	00:A0:A5:5D:2B:83	15	15	
0/16	00:A0:A5:5D:2B:83	16	16	
1/1	00:A0:A5:5D:2B:83	54	54	
1/2	00:A0:A5:5D:2B:83	55	55	
1/3	00:A0:A5:5D:2B:83	56	56	
1/4	00:A0:A5:5D:2B:83	57	57	
1/5	00:A0:A5:5D:2B:83	58	58	

Figure 2-44: Port Description

Table 2-43: Port Description Fields

Field	Description
Slot/Port	Select the interface for which data is to be displayed or configured.
Port Description	Enter text to describe a port. It can be up to 64 characters in length. The description can contain spaces and non-alphanumeric characters.
Slot/Port	Identifies the port.
Physical Address	Displays the physical address of the specified interface.
PortList Bit Offset	Displays the bit offset value which corresponds to the port when the MIB object type PortList is used to manage in SNMP.
IfIndex	Displays the interface index associated with the port.
Port Description	Shows the configured port description. By default, the port does not have an associated description.

- If you change a port description, click **Submit** to apply the change to the system.
- Click **Refresh** to redisplay the page with the latest information from the router.

2.9.4 Cable Test

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

Note...

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

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

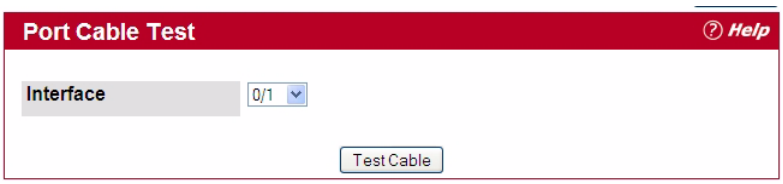

Figure 2-45: Cable Test

The page displays with additional fields when you click **Test Cable**.

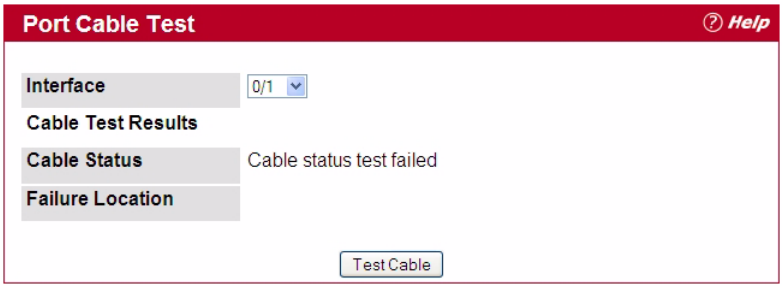

Figure 2-46: Cable Test

Table 2-44: Cable Test Fields

Field	Description
Slot/Port	This field indicates the interface to which the cable to be tested is connected.
Interface	Displays the interface tested in the Slot/Port notation.
Cable Status	This displays the cable status as Normal, Open, or Short. • **Normal**: The cable is working correctly. • **Open**: The cable is disconnected or there is a faulty connector. • **Short**: There is an electrical short in the cable. • **Cable Test Failed**: The cable status could not be determined. The cable may in fact be working. This field is displayed after you click Test Cable and results are available.
Cable Length	Displays the estimated length of the cable in meters. The length is displayed as a range between the shortest estimated length and the longest estimated length. Unknown is displayed if the cable length could not be determined. This field is displayed only when the cable status is Normal.
Failure Location	The estimated distance in meters from the end of the cable to the failure location. The failure location is only displayed if the cable status is Open or Short.

Select a port and click **Test Cable** to display its status.

If the port has an active link while the cable test is run, the link can go down for the duration of the test. The test may take several seconds to run.

The command returns a cable length estimate if this feature is supported by the PHY for the current link speed. Note that if the link is down and a cable is attached to a 10/100 Ethernet adapter, then the cable status may display as Open or Short because some Ethernet adapters leave unused wire pairs unterminated or grounded.

2.9.5 Multiple Port Mirroring

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

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

Use the Multiple Port Mirroring page to define port mirroring sessions.

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

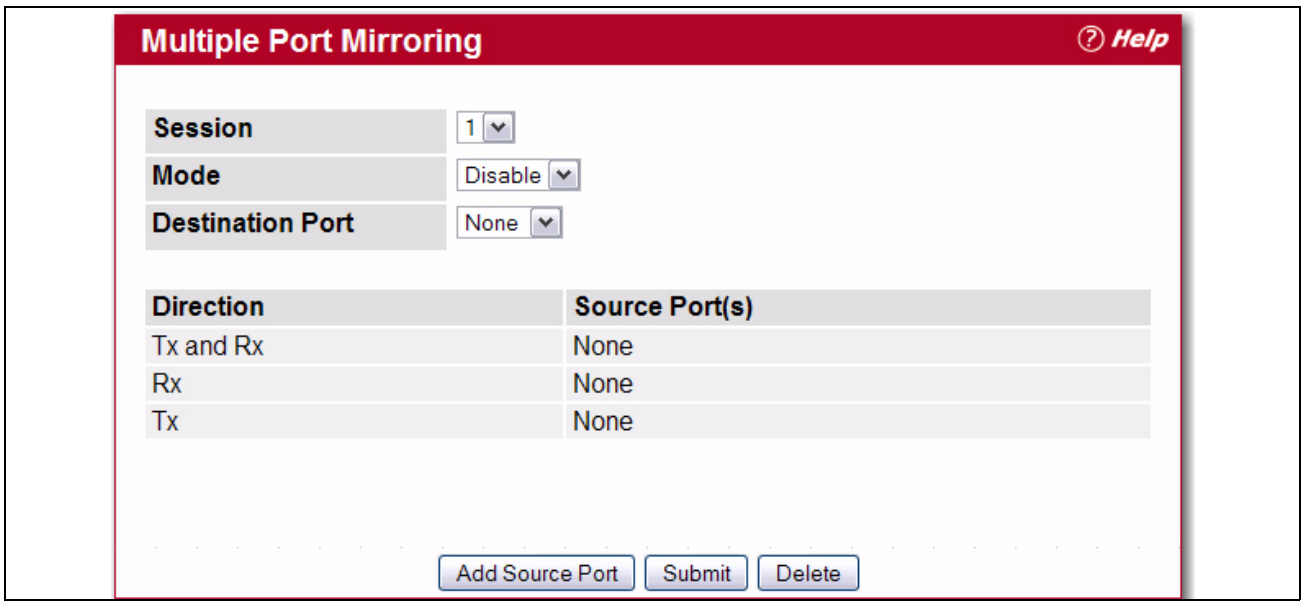

Figure 2-47: Multiple Port Mirroring

Table 2-45: Multiple Port Mirroring Fields

Field	Description
Session	Specifies the monitoring session.
Mode	Enables you to turn on of off Multiple Port Mirroring. The default is Disabled (off).
Destination Port	Select the port to which port traffic may be copied.
Direction	Specifies the direction of traffic on source port(s) which will be sent to the probe port. Possible values are: • Tx and Rx: Both Ingress and Egress traffic. • Rx: Ingress traffic only. • Tx: Egress traffic only.

2.9.5.1 Adding a Port Mirroring Session

Note...

A Port will be removed from a VLAN or LAG when it becomes a destination mirror.

1. From the Port Mirroring page, click **Add** to display the **Add Source Ports** page.

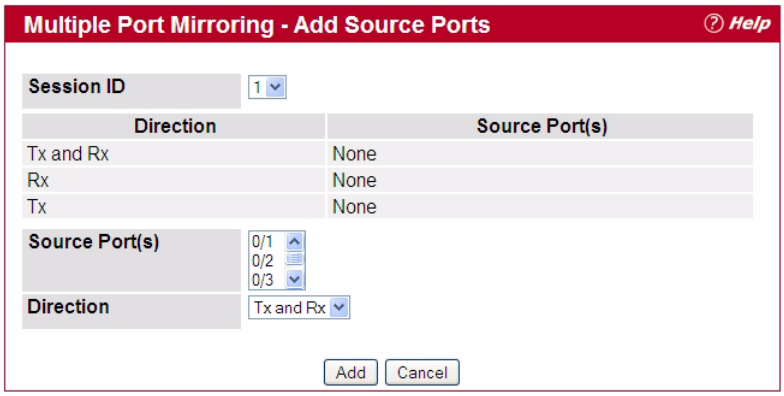

Figure 2-48: Multiple Port Mirroring—Add Source Ports

2. Configure the following fields:

Table 2-46: Multiple Port Mirroring—Add Source Ports Fields

Field	Description
Session	Specifies the monitoring session.
Source Port	Select the unit and port from which traffic is mirrored. Up to eight source ports can be mirrored to a destination port.
Direction	Select the type traffic monitored on the source port, which can be one of the following: • **Tx and Rx**: Monitors transmitted and received packets. • **Rx**: Monitors received packets only. • **Tx**: Monitors transmitted packets only.

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

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

2.9.5.2 Removing or Modifying a Port Mirroring Session

1. From the Port Mirroring page, click **Remove Source Port**.
2. Select one or more source ports to remove from the session.

   Use the CTRL key to select multiple ports to remove.
3. Click **Remove**.

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

2.10 TR-069 Client

TR-069 is a bidirectional remote management specification for customer premises equipment (CPE). TR-069 defines the CPE WAN Management Protocol (CWMP), which enables communication between the CPE and a remote auto-configuration server (ACS) to perform auto-configuration, dynamic service provisioning, software/firmware image management, status and performance monitoring, and diagnostics on the CPE.

TR-069 functions independently from the AutoInstall feature (see 2.14.13AutoInstall114). AutoInstall attempts to download a configuration file from a TFTP server if none is available locally when the switch boots. You do not have to enable TR-069 in order to use AutoInstall. However, if an AutoInstall attempt fails and TR-069 is configured on the CPE, then the CPE will attempt to obtain a configuration file from the ACS using TR-069.

The TR-069 client ACS URL can also be learned automatically during DHCP address assignment of the WAN port. This is possible by configuring the ACS URL as part of option 43: sub-option 1 in DHCP Server configuration.

2.10.1 TR-069 Configuration

Use the TR-069 configuration page to configure this feature on the switch. To display this page, click **System > TR-069 > TR-069 Configuration**.

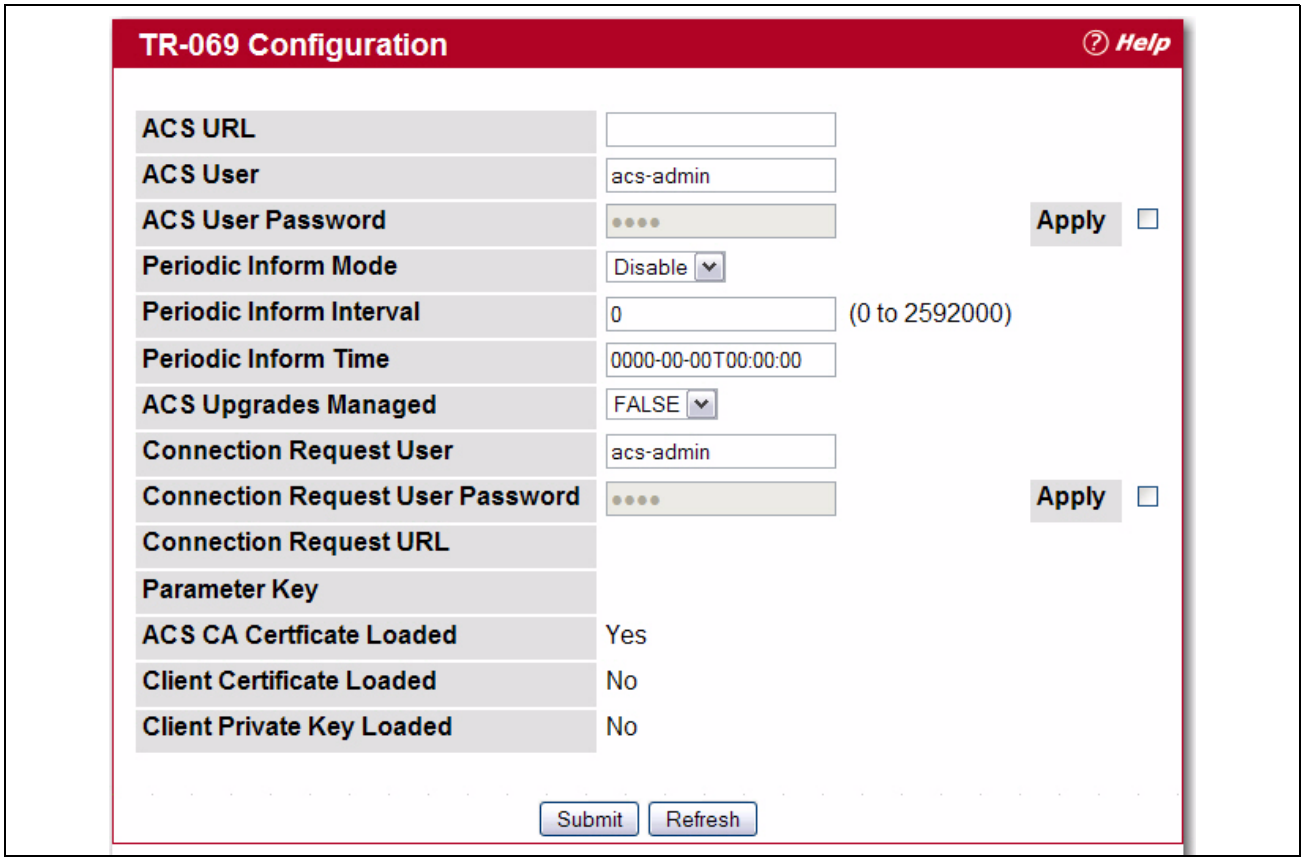

Figure 2-49: TR-069 Configuration

Table 2-47: TR-069 Configuration

Field	Description
ACS URL	Enter the URL for the CPE to connect to the ACS using the CPE WAN Management Protocol.
ACS User	Enter the user name for authenticating the CPE when it makes a TR-069 connection to the ACS. This parameter is used only when SSL support is not present.
ACS User Password	Enter the user password for authenticating the CPE when in makes a connection to the ACS.
Periodic Inform Mode	Indicates whether or not the CPE sends CPE information to the ACS using Periodic Inform Messages.
Periodic Inform Interval	Enter the duration in seconds of the interval in which the CPE attempts to connect with the ACS when Periodic Inform mode is enabled.
Periodic Inform Time	Enter the time when the CPE should initiate the inform messages. Each inform message must occur at this reference time plus or minus an integer multiple of the Periodic Inform Interval. A zero value (0000 0000T00:00:00) Indicates that no particular time reference is specified. That is, the CPE chooses the time reference but adheres to the specified Periodic Inform Interval.
ACS Upgrades Managed	Select whether or not the ACS will manage upgrades for the CPE. If **True**, the CPE cannot use the user interfaces (CLI, Web, and SNMP) for upgrades. If **False**, the CPE can use these interfaces to perform software upgrades.
Connection Request User	Enter the user name (up to 256 alphanumeric characters) for authenticating an ACS when it makes a connection request to the CPE.
Connection Request User Password	Enter the password (up to 256 alphanumeric characters) for authenticating an ACS when it makes a connection request to the CPE.
Connection Request URL	Enter the user HTTP URL for an ACS to make a connection request notification to the CPE.
Param Key	A max length 32 character string that provides a means to track the last successful transaction.
ACS CA Certificate Loaded	Indicates whether the ACS certification authority SSL certificate is present and loaded by the CPE.
Client Certificate Loaded	Indicates whether the CPE client SSL certificate is present and loaded by the CPE. This is used during certificate based CPE authentication.
Client Private Key Loaded	Indicates whether the CPE SSL private key is present and loaded by the CPE. This is used during certificate based CPE authentication.

If you make any changes to the page, click Submit to apply the changes to the system.

2.10.2 TR-069 Statistics

Use this page to view statistics on TR-069 communication with the ACS.

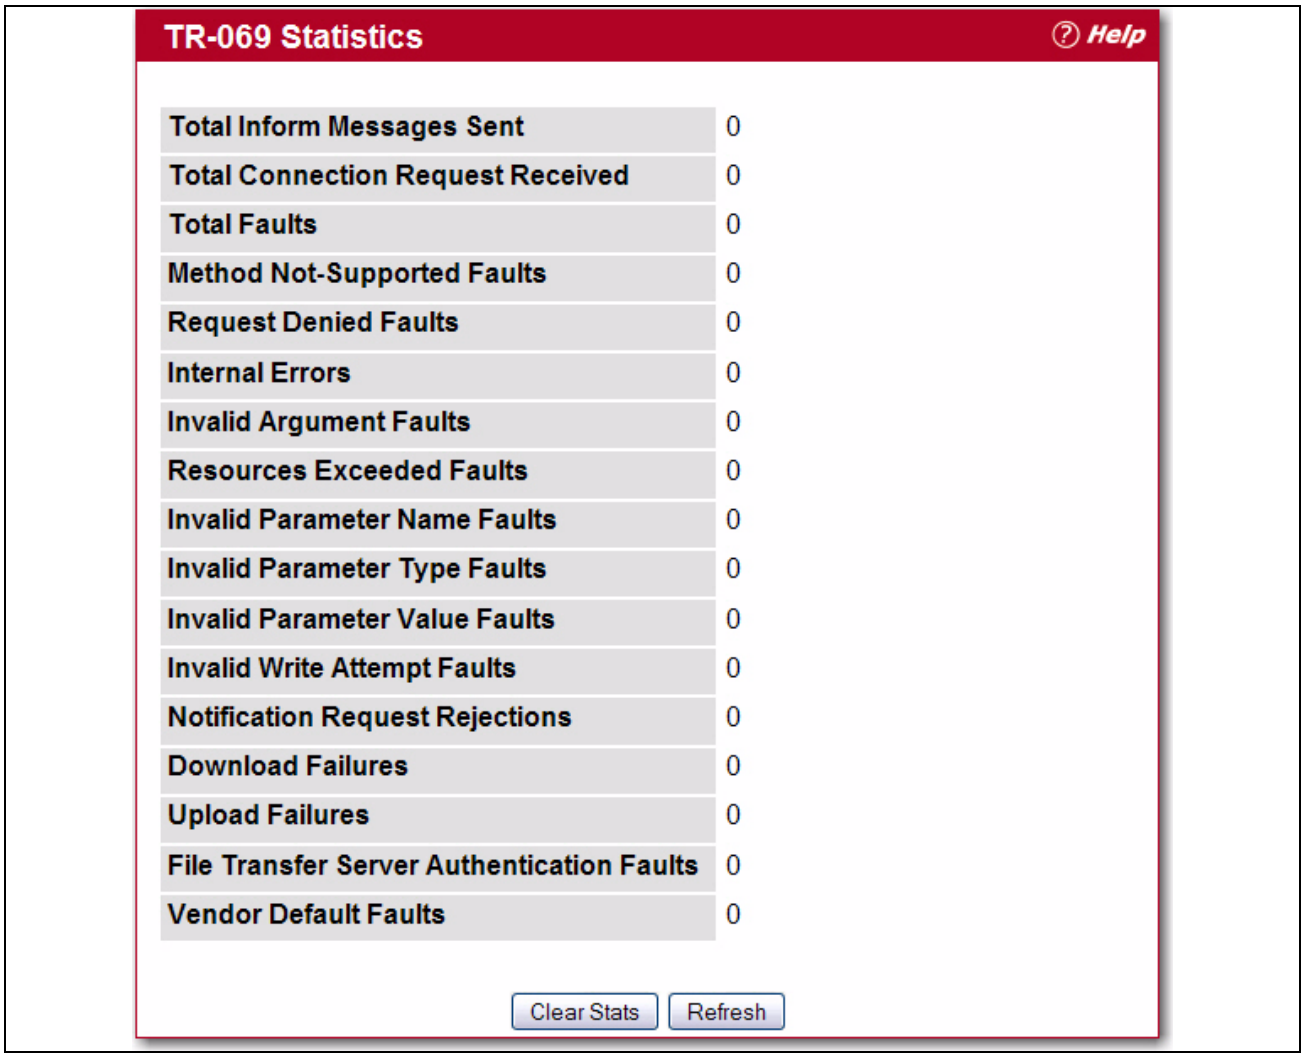

Figure 2-50: TR-069 Statistics

Table 2-48: TR-069 Statistics

Field	Description
Total Inform Messages Sent	Number of inform messages sent by the CPE since the last system reset.
Total Connection Requests Received	Number of connection request messages received by the CPE since the last system reset.
Total Faults	Number of faults encountered by the CPE since the last system reset.
Method Not-Supported Faults	Number of RPC requests with an unsupported RPC method received by the CPE since the last system reset.
Request Denied Faults	Number of RPC requests denied by the CPE since the last system reset.
Internal Errors	Number of RPC requests failed due to internal processing errors by the CPE since the last system reset.
Invalid Argument Faults	Number of RPC methods with invalid arguments received by the CPE since the last system reset.
Resources Exceeded Faults	Number of errors occurred due to unavailability of resources at the CPE since the last system reset.
Invalid Parameter Name Faults	Number of RPC methods with invalid parameter names received by the CPE since the last system reset.
Invalid Parameter Type Faults	Number of RPC methods with invalid parameter names received by the CPE since the last system reset.
Invalid Parameter Value Faults	Number of RPC methods with invalid parameter values received by the CPE since the last system reset.
Invalid Write Attempt Faults	Number of attempts to set a non writable parameter by the CPE since the last system reset.
Notification Request Rejections	Number of RPC methods denied by the CPE since the last system reset.
Download Failures	Number of download failures encountered by the CPE since the last system reset.
Upload Failures	Number of upload failures encountered by the CPE since the last system reset.
File Transfer Server Authentication Failures	Number of file server authentication failures encountered by the CPE since the last system reset.
Vendor Default Faults	Number of vendor-defined errors encountered by the CPE since the last system reset.

Click **Clear Stats** to reset all counters to their initial values.

Click **Refresh** to display the latest cumulative data from the switch.

2.11 Configuring sFlow

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

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

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

2.11.1 sFlow Agent Summary

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

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

To access the sFlow Agent Summary page, click **System > sFlow > Agent Summary** in the navigation tree.

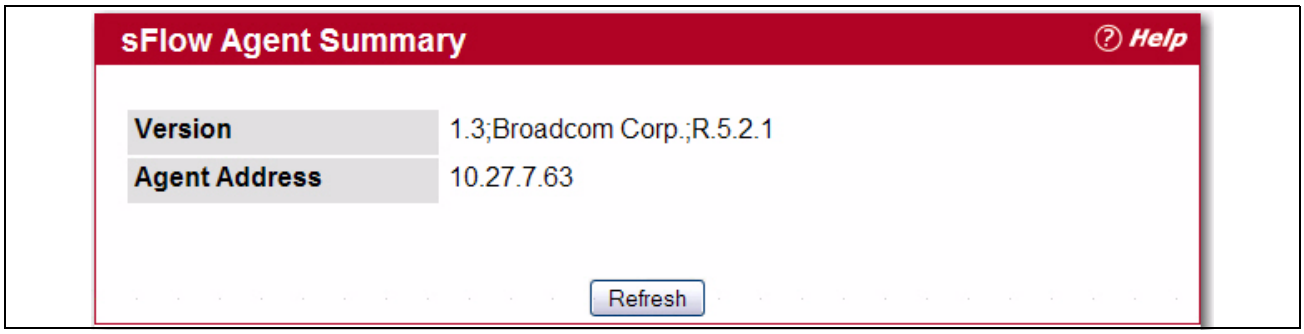

Figure 2-51: sFlow Agent Summary

Table 2-49: sFlow Agent Summary

Field	Description
Version	Uniquely identifies the version and implementation of this MIB. The version string must have the following structure: MIB Version;Organization;Software Revision where: • MIB Version: '1.3', the version of this MIB. • Organization: Broadcom Corp. • Revision: 1.0
Agent Address	The IP address associated with this agent.

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

2.11.2 sFlow Receiver Configuration

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

To access the sFlow Receiver Configuration page, click **System > sFlow > Receiver Configuration** in the navigation tree.

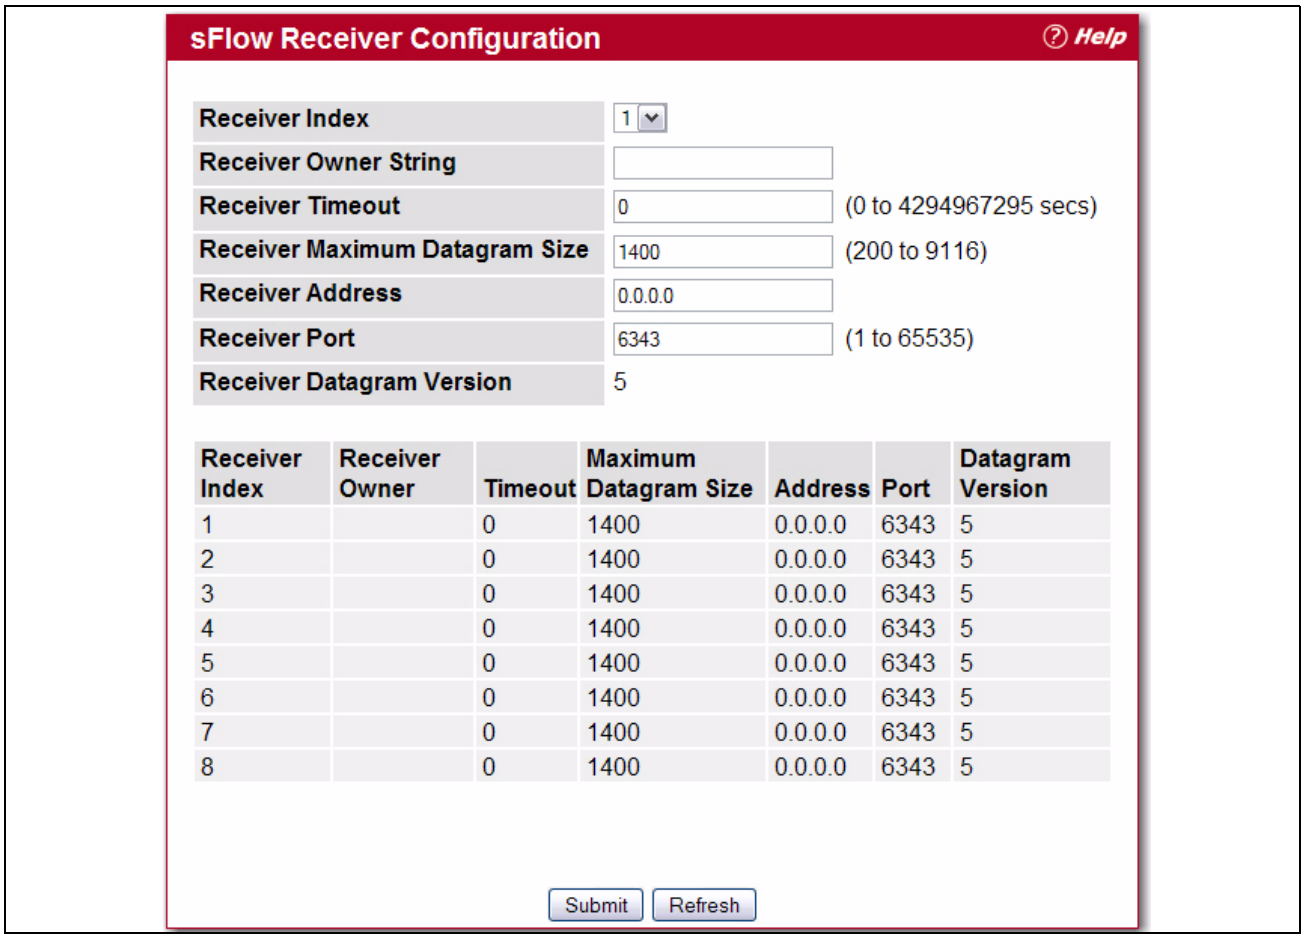

Figure 2-52: sFlow Receiver Configuration

Table 2-50: sFlow Receiver Configuration

Field	Description
Receiver Index	Selects the receiver for which data is to be displayed or configured. The allowed range is 1 to 8.
Receiver Owner String	The entity making use of this sFlowRcvrTable entry. The empty string indicates that the entry is currently unclaimed and the receiver configuration is reset to the default values. An entity wishing to claim an sFlowRcvrTable entry must ensure that the entry is unclaimed before trying to claim it. The entry is claimed by setting the owner string. The entry must be claimed before any changes can be made to other sampler objects.
sFlow Receiver Timeout	The time (in seconds) remaining before the sampler is released and stops sampling. A management entity wanting to maintain control of the sampler is responsible for setting a new value before the old one expires. The allowed range is 0 to 4294967295 seconds. A value of zero sets the selected receiver configuration to its default values.
sFlow Receiver Maximum Datagram Size	The maximum number of data bytes that can be sent in a single sample datagram. The manager should set this value to avoid fragmentation of the sFlow datagrams. The default value is 1400. The allowed range is 200 to 9116.)
sFlow Receiver Address	The IP address of the sFlow collector. If set to 0.0.0.0 no sFlow datagrams will be sent.
sFlow Receiver Port	The destination port for sFlow datagrams. The allowed range is 1 to 65535).
Receiver Datagram Version	The version of sFlow datagrams that should be sent.

- Use the **Submit** button to sent updated data to the switch and cause the changes to take effect on the switch.
- Use the **Refresh** button to refresh the page with the most current data from the switch.

2.11.3 sFlow Poller Configuration

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

2.11.3.1 Counter Sampling

The primary objective of Counter Sampling is to efficiently, periodically export counters associated with Data Sources. A maximum Sampling Interval is assigned to each sFlow instance associated with a Data Source.

Counter Sampling is accomplished as follows:

> The sFlow Agent keeps a list of counter sources being sampled. When a Packet Flow Sample is generated, the sFlow Agent examines the list and adds counters to the sample datagram, least recently sampled first. Counters are only added to the datagram if the sources are within a short period, i.e. five seconds, of failing to meet the required Sampling Interval. Periodically, i.e. every second, the sFlow Agent examines the list of counter sources and sends any counters that need to be sent to meet the sampling interval requirement.

To access the sFlow Poller Configuration page, click **System > sFlow > Poller Configuration** in the navigation tree.

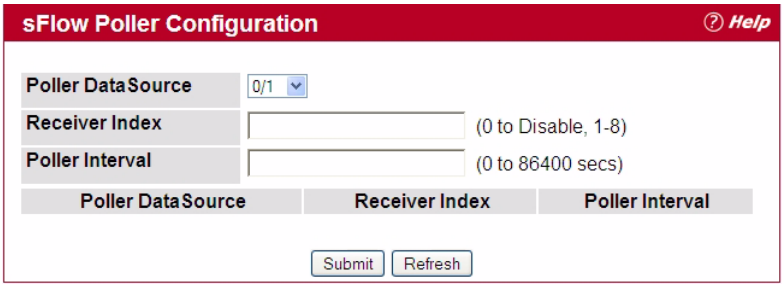

Figure 2-53: sFlow Poller Configuration

Table 2-51: sFlow Poller Configuration

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

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

2.11.4 sFlow Sampler Configuration

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

2.11.4.1 Packet Flow Sampling

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

Packet Flow Sampling is accomplished as follows:

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

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

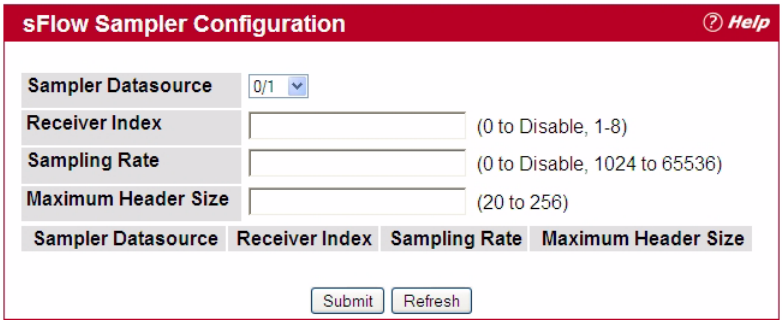

Figure 2-54: sFlow Sampler Configuration

Table 2-52: sFlow Sampler Configuration

Field	Description
Poller DataSource	The sFlow Datasource for this sFlow sampler. This Agent will support Physical ports only.
Receiver Index	The sFlow Receiver for this sFlow sampler. If set to zero, no packets will be sampled. Only active receivers can be set. If a receiver expires, then all samplers associated with the receiver will also expire. The allowed range is 1 to 8.
Sampling Rate	The statistical sampling rate for packet sampling from this source. A sampling rate of one (1) counts all packets. A sampling rate of zero (0) disables sampling. The allowed range is 1024 to 65536.
Maximum Header Size	The maximum number of bytes that should be copied from a sampled packet. The allowed range is 20 to 256.

2.12 Defining SNMP Parameters

Simple Network Management Protocol (SNMP) provides a method for managing network devices. The device supports SNMP version 1, SNMP version 2, and SNMP version 3. The Web interfaces supports configuration of SNMPv1 and v2; SNMPv3 is supported only in the CLI.

2.12.1 SNMP v1 and v2

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

2.12.2 SNMP v3

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

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

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

- Security
- Feature Access Control
- Traps

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

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

2.12.3 SNMP Community Configuration

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

Use the Community Configuration page to enable SNMP and Authentication notifications.

To display the Community Configuration page, click **System > SNMP > Community Configuration** in the navigation tree.

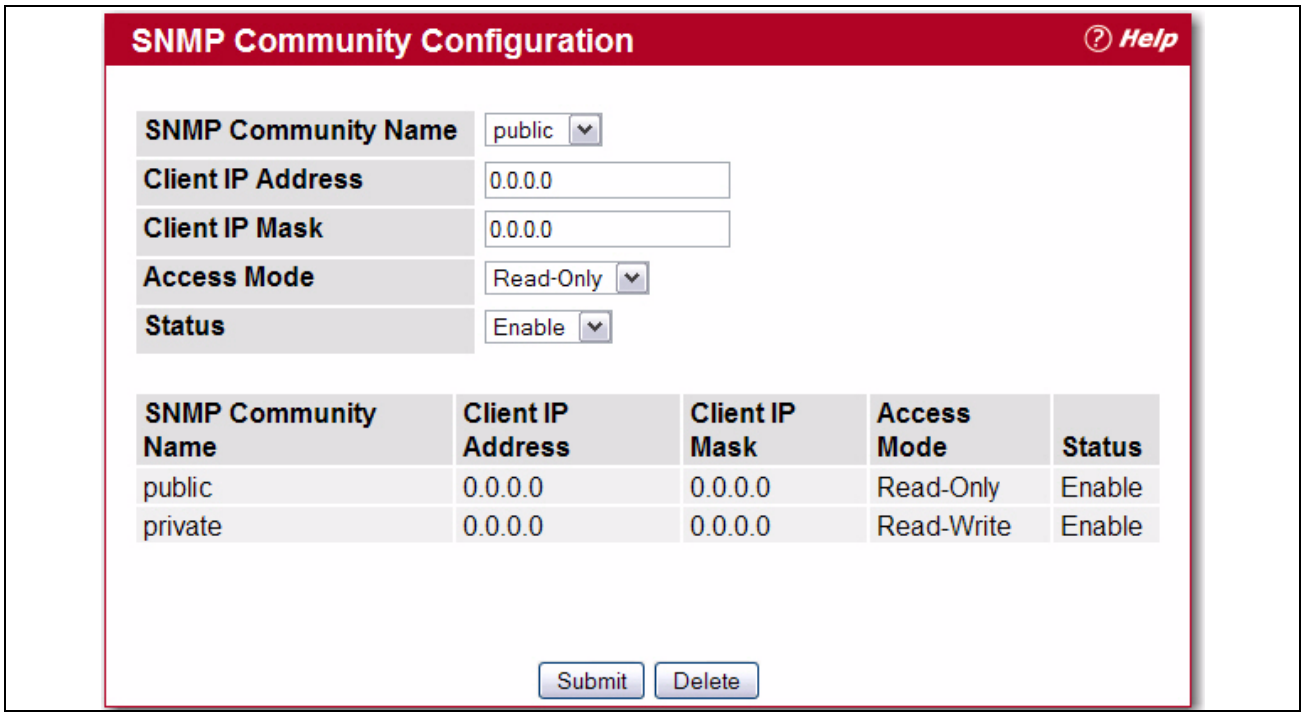

Figure 2-55: SNMP Community Configuration

Table 2-53: Community Configuration Fields

Field	Description
Community	Contains the predefined and user-defined community strings that act as a password and are used to authenticate the SNMP management station to the device. A community string can contain a maximum of 20 characters. By default, the options available in the menu are as follows: • **public**: This SNMP community has Read Only privileges and its status set to enable • **private**: This SNMP community has Read/Write privileges and its status set to enable.
SNMP Community Name	Use this field to reconfigure an existing community or to create a new one. A valid entry is a case-sensitive string of up to 16 characters.

Table 2-53: Community Configuration Fields (Continued)

Field	Description
Client IP Address	Taken together, the **Client IP Address** and **Client IP Mask** denote a range of IP addresses from which SNMP clients may use that community to access this device. If either (IP Address or IP Mask) value is 0.0.0.0, access is allowed from any IP address. Otherwise, every client's IP address is ANDed with the mask, as is the Client IP Address, and, if the values are equal, access is allowed. For example, if the Client IP Address and Client IP Mask parameters are 192.168.1.0/255.255.255.0, then any client whose IP address is 192.168.1.0 through 192.168.1.255 (inclusive) will be allowed access. To allow access from only one station, use a Client IP Mask value of 255.255.255.255, and use that machine's IP address for Client IP Address.
Client IP Mask	Along with the **Client IP Address**, the **Client IP Mask** denotes a range of IP addresses from which SNMP clients may use that community to access this device.
Access Mode	Specify the access level for this community: • **Read-Only**: The Community has read only access to the MIB objects configured in the view. • **Read-Write**: The Community has read/modify access to the MIB objects configured in the view.
Status	Specify the status of this community: • **Enable**: The community is enabled, and the Community Name must be unique among all valid Community Names or the set request will be rejected. • **Disable**: The Community is disabled and the Community Name becomes invalid.

- If you make any changes to the page, click **Submit** to apply the changes to the system. If you create a new Community, it is added to the table below the **Submit** button.
- Click **Delete** to delete the selected SNMP Community.

2.12.4 Trap Receiver Configuration

Use the Trap Receiver Configuration page to configure information about the SNMP community and the trap manager that will receive its trap packets.

To access the Trap Receiver Configuration page, click **System > SNMP > Trap Receiver Configuration** from the navigation tree.

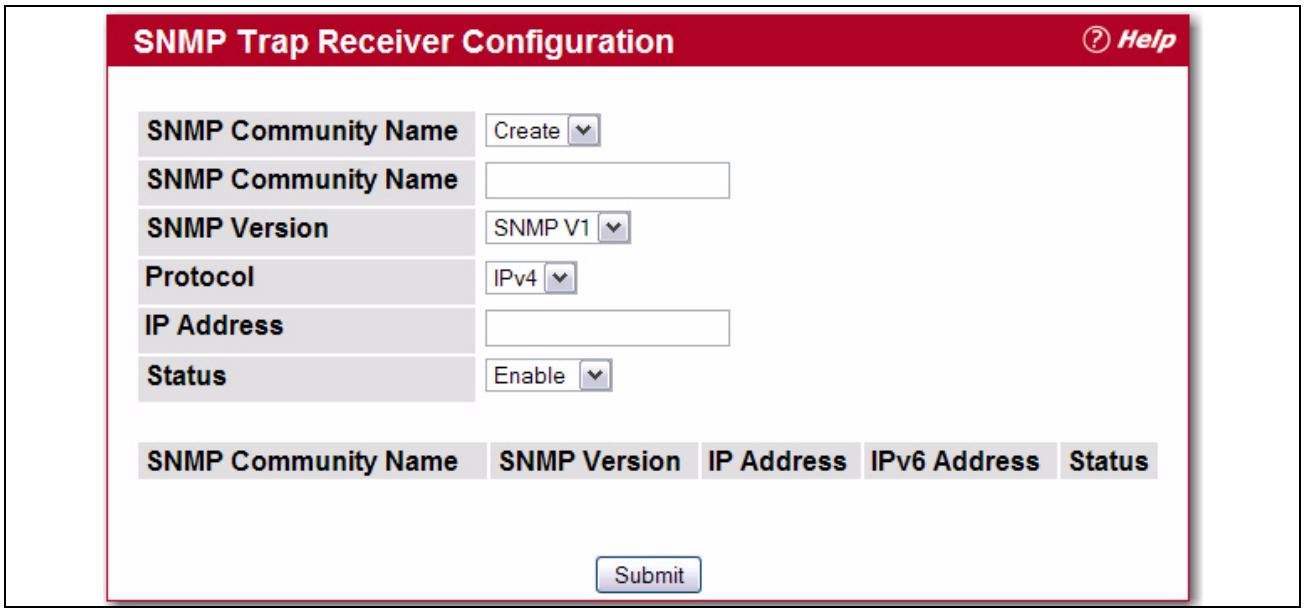

Figure 2-56: Trap Receiver Configuration

Table 2-54: Trap Receiver Configuration Fields

Field	Description
(Create) SNMP Community Name	When this field is set to **Create**, you can configure new SNMP trap receiver information in the rest of the fields. If you have already configured an SNMP trap receiver, you can select it from the drop-down menu to change the settings or delete it.
SNMP Community Name	Enter the community string for the SNMP trap packet to be sent to the trap manager. This may be up to 16 characters and is case sensitive.
SNMP Version	Select the trap version to be used by the receiver from the pull down menu: • **SNMP v1**. Uses SNMP v1 to send traps to the receiver. • **SNMP v2**. Uses SNMP v2 to send traps to the receiver.
Protocol	Select the type of protocol used for the SNMP Trap Receiver Configuration: • **IPv4**. Choose IPv4 to enter the address in IPv4 format. • **IPv6**. Choose IPv6 to enter the address in IPv6 format.
IP Address	Enter the IP address in dotted-decimal format to receive SNMP traps from this device.
Status	Select the receiver's status from the pulldown menu: • **Enable**: Send traps to the receiver • **Disable**: Do not send traps to the receiver.

If you make any changes to the page, click **Submit** to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must perform a **save**.

2.12.5 Supported MIBs

The Supported MIBs page lists the MIBs that the system currently supports.

To access the Supported MIBs page, click **System > SNMP > Supported MIBs** in the navigation menu. A portion of the web screen is shown Figure 2-57.

SNMP Supported MIBs Help

Name	Description
RFC 1907 - SNMPv2-MIB	The MIB module for SNMPv2 entities
RFC 2819 - RMON-MIB	Remote Network Monitoring Management Information Base
Broadcom-REF-MIB	Broadcom Reference
SNMP-COMMUNITY-MIB	This MIB module defines objects to help support coexistence between SNMPv1, SNMPv2, and SNMPv3.
SNMP-FRAMEWORK-MIB	The SNMP Management Architecture MIB
SNMP-MPD-MIB	The MIB for Message Processing and Dispatching
SNMP-NOTIFICATION-MIB	The Notification MIB Module
SNMP-TARGET-MIB	The Target MIB Module

Figure 2-57: Supported MIBs

Table 2-55: Supported MIBs Fields

Field	Description
Name	The RFC number if applicable and the name of the MIB.
Description	The RFC title or MIB description.

2.13 Viewing System Statistics

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

2.13.1 Switch Detailed

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

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

Switch Detailed Statistics Help

Interface	417
Octets Received	0
Packets Received Without Error	0
Unicast Packets Received	0
Multicast Packets Received	0
Broadcast Packets Received	0
Receive Packets Discarded	0
Octets Transmitted	492
Packets Transmitted Without Errors	6
Unicast Packets Transmitted	0
Multicast Packets Transmitted	6
Broadcast Packets Transmitted	0
Transmit Packets Discarded	0
Most Address Entries Ever Used	1
Address Entries in Use	1
Maximum VLAN Entries	4093
Most VLAN Entries Ever Used	2
Static VLAN Entries	2
Dynamic VLAN Entries	0
VLAN Deletes	0
Time Since Counters Last Cleared	0 day 20 hr 42 min 26 sec

Clear Counters Refresh

Figure 2-58: Switch Detailed

Table 2-56: Switch Detailed Statistics Fields

Field	Description
fIndex	This object indicates the ifIndex of the interface table entry associated with the processor of this switch.
Octets Received	The total number of octets of data received by the processor (excluding framing bits but including FCS octets).
Unicast Packets Received	The number of subnetwork-unicast packets delivered to a higher-layer protocol.
Multicast Packets Received	The total number of packets received that were directed to a multicast address. Note that this number does not include packets directed to the broadcast address.
Broadcast Packets Received	The total number of packets received that were directed to the broadcast address. Note that this does not include multicast packets.
Receive Packets Discarded	The number of inbound packets which were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher-layer protocol. A possible reason for discarding a packet could be to free up buffer space.
Octets Transmitted	The total number of octets transmitted out of the interface, including framing characters.
Packets Transmitted Without Errors	The total number of packets transmitted out of the interface.
Unicast Packets Transmitted	The total number of packets that higher-level protocols requested be transmitted to a subnetwork-unicast address, including those that were discarded or not sent.
Multicast Packets Transmitted	The total number of packets that higher-level protocols requested be transmitted to a Multicast address, including those that were discarded or not sent.
Broadcast Packets Transmitted	The total number of packets that higher-level protocols requested be transmitted to the Broadcast address, including those that were discarded or not sent.
Transmit Packets Discarded	The number of outbound packets which were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher-layer protocol. A possible reason for discarding a packet could be to free up buffer space.
Most Address Entries Ever Used	The highest number of Forwarding Database Address Table entries that have been learned by this switch since the most recent reboot.
Address Entries in Use	The number of Learned and static entries in the Forwarding Database Address Table for this switch.
Maximum VLAN Entries	The maximum number of Virtual LANs (VLANs) allowed on this switch.
Most VLAN Entries Ever Used	The largest number of VLANs that have been active on this switch since the last reboot.
Static VLAN Entries	The number of presently active VLAN entries on this switch that have been created statically.
Dynamic VLAN Entries	The number of presently active VLAN entries on this switch that have been created by GVRP registration.
VLAN Deletes	The number of VLANs on this switch that have been created and then deleted since the last reboot.
Time Since Counters Last Cleared	The elapsed time, in days, hours, minutes, and seconds, since the statistics for this switch were last cleared.

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

2.13.2 Switch Summary

Use the Switch Summary page to view a summary of statistics for traffic on the switch.

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

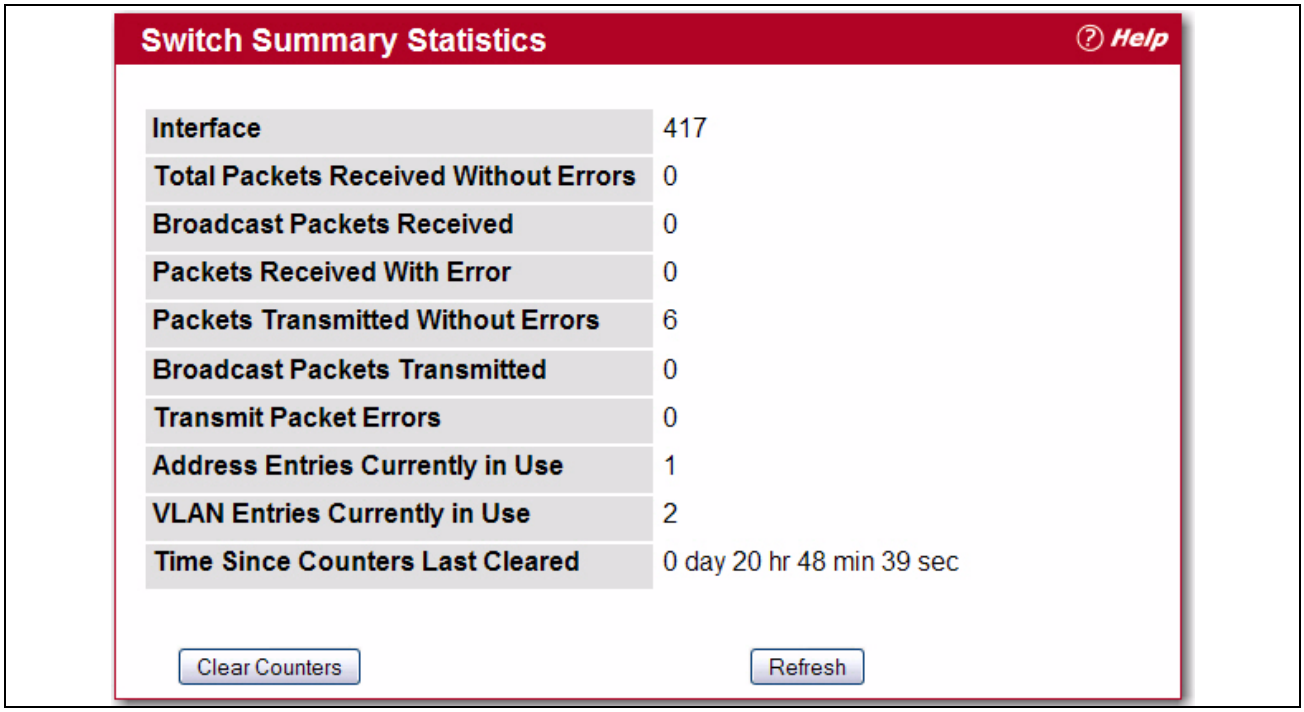

Figure 2-59: Switch Summary

Table 2-57: Switch Summary Fields

Field	Description
ifIndex	This object indicates the ifIndex of the interface table entry associated with the Processor of this switch.
Total Packets Received Without Errors	The total number of packets, including multicast packets, that were directed to the broadcast address.
Broadcast Packets Received	The total number of packets received that were directed to the broadcast address. Note that this does not include multicast packets.
Packets Received With Error	The number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol.
Packets Transmitted Without Errors	The total number of packets transmitted out of the interface.
Broadcast Packets Transmitted	The total number of packets that higher-level protocols requested to be transmitted to the Broadcast address, including those that were discarded or not sent.
Transmit Packet Errors	The number of outbound packets that could not be transmitted because of errors.
Address Entries Currently in Use	The total number of Forwarding Database Address Table entries now active on the switch, including learned and static entries.
VLAN Entries Currently in Use	The number of VLAN entries presently occupying the VLAN table.
Time Since Counters Last Cleared	The elapsed time, in days, hours, minutes, and seconds since the statistics for this switch were last cleared.

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

2.13.3 Port Detailed

The Port Detailed page displays a variety of per-port traffic statistics.

To access the Port Detailed page, click **System > Statistics > Port Detailed** in the navigation tree.

Figure 2-60 shows some, but not all, of the fields on the Port Detailed page.

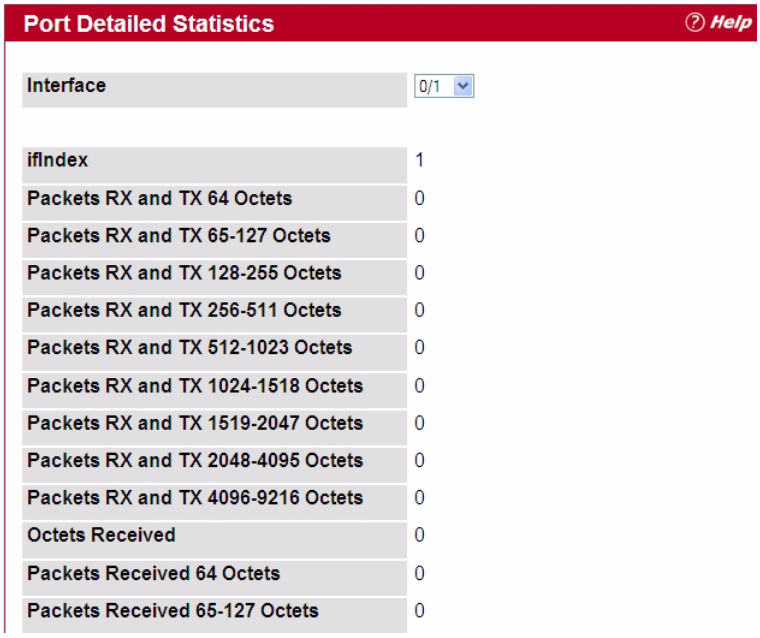

Figure 2-60: Port Detailed

Table 2-58: Port Fields

Field	Description
Slot/Port	Use the drop-down menu to select the interface for which data is to be displayed or configured.
ifIndex	This field indicates the ifIndex of the interface table entry associated with this port on an adapter.
Packets RX and TX 64 Octets	The total number of packets (including bad packets) received or transmitted that were 64 octets in length (excluding framing bits but including FCS octets).
Packets RX and TX 65-127 Octets	The total number of packets (including bad packets) received or transmitted that were between 65 and 127 octets in length inclusive (excluding framing bits but including FCS octets).
Packets RX and TX 128-255 Octets	The total number of packets (including bad packets) received or transmitted that were between 128 and 255 octets in length inclusive (excluding framing bits but including FCS octets).
Packets RX and TX 256-511 Octets	The total number of packets (including bad packets) received or transmitted that were between 256 and 511 octets in length inclusive (excluding framing bits but including FCS octets).

Table 2-58: Port Fields (Continued)

Field	Description
Packets RX and TX 512-1023 Octets	The total number of packets (including bad packets) received or transmitted that were between 512 and 1023 octets in length inclusive (excluding framing bits but including FCS octets).
Packets RX and TX 1024-1518 Octets	The total number of packets (including bad packets) received or transmitted that were between 1024 and 1518 octets in length inclusive (excluding framing bits but including FCS octets).
Packets RX and TX 1519-1522 Octets	The total number of packets (including bad packets) received or transmitted that were between 1519 and 1522 octets in length inclusive (excluding framing bits but including FCS octets).
Packets RX and TX 1523-2047 Octets	The total number of packets (including bad packets) received or transmitted that were between 1523 and 2047 octets in length inclusive (excluding framing bits but including FCS octets).
Packets RX and TX 2048-4095 Octets	The total number of packets (including bad packets) received or transmitted that were between 2048 and 4095 octets in length inclusive (excluding framing bits but including FCS octets).
Packets RX and TX 4096-9216 Octets	The total number of packets (including bad packets) received or transmitted that were between 4096 and 9216 octets in length inclusive (excluding framing bits but including FCS octets).
Octets Received	The total number of octets of data (including those in bad packets) received on the network (excluding framing bits but including FCS octets). This object can be used as a reasonable estimate of ethernet utilization. If greater precision is desired, the etherStatsPkts and etherStatsOctets objects should be sampled before and after a common interval.
Packets Received 64 Octets	The total number of packets (including bad packets) received that were 64 octets in length (excluding framing bits but including FCS octets).
Packets Received 65-127 Octets	The total number of packets (including bad packets) received that were between 65 and 127 octets in length inclusive (excluding framing bits but including FCS octets).
Packets Received 128-255 Octets	The total number of packets (including bad packets) received that were between 128 and 255 octets in length inclusive (excluding framing bits but including FCS octets).
Packets Received 256-511 Octets	The total number of packets (including bad packets) received that were between 256 and 511 octets in length inclusive (excluding framing bits but including FCS octets).
Packets Received 512-1023 Octets	The total number of packets (including bad packets) received that were between 512 and 1023 octets in length inclusive (excluding framing bits but including FCS octets).
Packets Received 1024-1518 Octets	The total number of packets (including bad packets) received that were between 1024 and 1518 octets in length inclusive (excluding framing bits but including FCS octets).
Packets Received > 1522 Octets	The total number of packets received that were longer than 1522 octets (excluding framing bits, but including FCS octets) and were otherwise well formed.
Total Packets Received Without Errors	The total number of packets received that were without errors.
Unicast Packets Received	The number of subnetwork-unicast packets delivered to a higher-layer protocol.
Multicast Packets Received	The total number of good packets received that were directed to a multicast address. Note that this number does not include packets directed to the broadcast address.
Broadcast Packets Received	The total number of good packets received that were directed to the broadcast address. Note that this does not include multicast packets.
Total Packets Received with MAC Errors	The total number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol.

Table 2-58: Port Fields (Continued)

Field	Description
Jabbers Received	The total number of packets received that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error). Note that this definition of jabber is different than the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition where any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms.
Fragments Received	The total number of packets received that were less than 64 octets in length with ERROR CRC (excluding framing bits but including FCS octets).
Undersize Received	The total number of packets received that were less than 64 octets in length with GOOD CRC (excluding framing bits but including FCS octets).
Alignment Errors	The total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, but had a bad Frame Check Sequence (FCS) with a non-integral number of octets.
Rx FCS Errors	The total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, but had a bad Frame Check Sequence (FCS) with an integral number of octets
Overruns	The total number of frames discarded as this port was overloaded with incoming packets, and could not keep up with the inflow.
Total Received Packets Not Forwarded	A count of valid frames received which were discarded (i.e., filtered) by the forwarding process.
Local Traffic Frames	The total number of frames dropped in the forwarding process because the destination address was located off of this port.
802.3x Pause Frames Received	A count of MAC Control frames received on this interface with an opcode indicating the PAUSE operation. This counter does not increment when the interface is operating in half-duplex mode.
Unacceptable Frame Type	The number of frames discarded from this port due to being an unacceptable frame type.
Multicast Tree Viable Discards	The number of frames discarded when a lookup in the multicast tree for a VLAN occurs while that tree is being modified.
Reserved Address Discards	The number of frames discarded that are destined to an IEEE 802.1 reserved address and are not supported by the system.
Broadcast Storm Recovery	The number of frames discarded that are destined for FF:FF:FF:FF:FF:FF when Broadcast Storm Recovery is enabled.
CFI Discards	The number of frames discarded that have CFI bit set and the addresses in RIF are in non-canonical format.
Upstream Threshold	The number of frames discarded due to lack of cell descriptors available for that packet's priority level.
Total Packets Transmitted (Octets)	The total number of octets of data (including those in bad packets) transmitted on the network (excluding framing bits but including FCS octets). This object can be used as a reasonable estimate of ethernet utilization. If greater precision is desired, the etherStatsPkts and etherStatsOctets objects should be sampled before and after a common interval.
Packets Transmitted 64 Octets	The total number of packets (including bad packets) received that were 64 octets in length (excluding framing bits but including FCS octets).
Packets Transmitted 65-127 Octets	The total number of packets (including bad packets) received that were between 65 and 127 octets in length inclusive (excluding framing bits but including FCS octets).
Packets Transmitted 128-255 Octets	The total number of packets (including bad packets) received that were between 128 and 255 octets in length inclusive (excluding framing bits but including FCS octets).
Packets Transmitted 256-511 Octets	The total number of packets (including bad packets) received that were between 256 and 511 octets in length inclusive (excluding framing bits but including FCS octets).

Table 2-58: Port Fields (Continued)

Field	Description
Packets Transmitted 512-1023 Octets	The total number of packets (including bad packets) received that were between 512 and 1023 octets in length inclusive (excluding framing bits but including FCS octets).
Packets Transmitted 1024-1518 Octets	The total number of packets (including bad packets) received that were between 1024 and 1518 octets in length inclusive (excluding framing bits but including FCS octets).
Packets Transmitted 1523-2047 Octets	The total number of packets (including bad packets) received that were between 1523 and 2047 octets in length inclusive (excluding framing bits but including FCS octets).
Packets Transmitted 2048-4095 Octets	The total number of packets (including bad packets) received that were between 2048 and 4095 octets in length inclusive (excluding framing bits but including FCS octets).
Packets Transmitted 4096-9216 Octets	The total number of packets (including bad packets) received that were between 4096 and 9216 octets in length inclusive (excluding framing bits but including FCS octets).
Maximum Frame Size	The maximum ethernet frame size the interface supports or is configured, including ethernet header, CRC, and payload. (1518 to 9216). The default maximum frame size is 1518.
Total Packets Transmitted Successfully	The number of frames that have been transmitted by this port to its segment.
Unicast Packets Transmitted	The total number of packets that higher-level protocols requested be transmitted to a subnetwork-unicast address, including those that were discarded or not sent.
Multicast Packets Transmitted	The total number of packets that higher-level protocols requested be transmitted to a Multicast address, including those that were discarded or not sent.
Broadcast Packets Transmitted	The total number of packets that higher-level protocols requested be transmitted to the Broadcast address, including those that were discarded or not sent.
Total Transmit Errors	The sum of Single, Multiple, and Excessive Collisions.
Tx FCS Errors	The total number of packets transmitted that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, but had a bad Frame Check Sequence (FCS) with an integral number of octets
Tx Oversized	The total number of frames that exceeded the max permitted frame size. This counter has a max increment rate of 815 counts per second at 10 Mb/s.
Underrun Errors	The total number of frames discarded because the transmit FIFO buffer became empty during frame transmission.
Total Transmit Packets Discarded	The sum of single collision frames discarded, multiple collision frames discarded, and excessive frames discarded.
Single Collision Frames	A count of the number of successfully transmitted frames on a particular interface for which transmission is inhibited by exactly one collision.
Multiple Collision Frames	A count of the number of successfully transmitted frames on a particular interface for which transmission is inhibited by more than one collision.
Excessive Collision Frames	A count of frames for which transmission on a particular interface fails due to excessive collisions.
Port Membership Discards	The number of frames discarded on egress for this port due to egress filtering being enabled.
STP BPDUs Received	Number of STP BPDUs received at the selected port.
STP BPDUs Transmitted	Number of STP BPDUs transmitted from the selected port.
RSTP BPDUs Received	Number of RSTP BPDUs received at the selected port.
RSTP BPDUs Transmitted	Number of RSTP BPDUs transmitted from the selected port.
MSTP BPDUs Received	Number of MSTP BPDUs received at the selected port.
MSTP BPDUs Transmitted	Number of MSTP BPDUs transmitted from the selected port.
802.3x Pause Frames Transmitted	A count of MAC Control frames transmitted on this interface with an opcode indicating the PAUSE operation. This counter does not increment when the interface is operating in half-duplex mode.

Table 2-58: Port Fields (Continued)

Field	Description
GVRP PDUs Received	The count of GVRP PDUs received in the GARP layer.
GVRP PDUs Transmitted	The count of GVRP PDUs transmitted from the GARP layer.
GVRP Failed Registrations	The number of times attempted GVRP registrations could not be completed.
GMRP PDUs Received	The count of GMRP PDUs received from the GARP layer.
GMRP PDUs Transmitted	The count of GMRP PDUs transmitted from the GARP layer.
GMRP Failed Registrations	The number of times attempted GMRP registrations could not be completed.
Time Since Counters Last Cleared	The elapsed time, in days, hours, minutes, and seconds since the statistics for this port were last cleared.

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

2.13.4 Port Summary

The Port Summary page shows a summary of per-port traffic statistics on the switch.

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

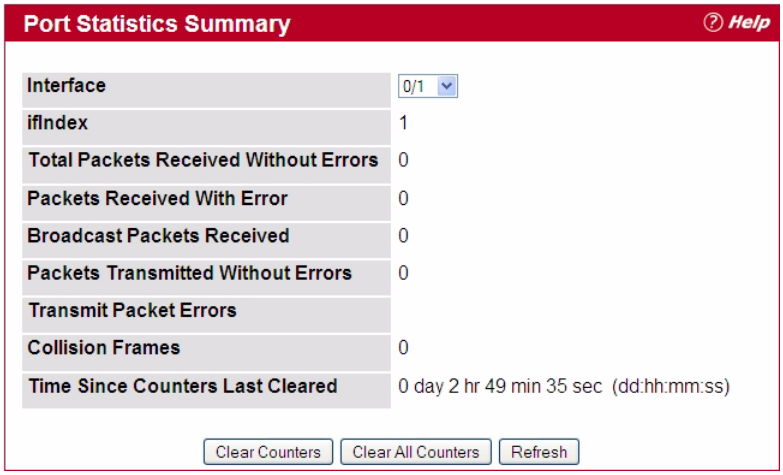

Figure 2-61: Port Summary

Table 2-59: Port Summary Fields

Field	Description
Slot/Port	Use the drop-down menu to select the interface for which data is to be displayed or configured.
ifIndex	This field indicates the ifIndex of the interface table entry associated with this port on an adapter.
Total Packets Received Without Errors	The total number of packets received that were without errors.
Packets Received With Error	The number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol.
Broadcast Packets Received	The total number of good packets received that were directed to the broadcast address. Note that this does not include multicast packets.
Packets Transmitted Without Errors	The number of frames that have been transmitted by this port to its segment.
Transmit Packet Errors	The number of outbound packets that could not be transmitted because of errors.
Collision Frames	The best estimate of the total number of collisions on this Ethernet segment.
Time Since Counters Last Cleared	The elapsed time, in days, hours, minutes, and seconds since the statistics for this port were last cleared.

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

2.14 Using System Utilities

The System Utilities folder contains links to the following Web pages that help you manage the switch:

- Save All Applied Changes
- System Reset
- Reset Configuration to Defaults
- Erase Startup Config File
- Reset Passwords to Defaults
- Download File To Switch (TFTP)
- Upload File From Switch (TFTP)
- Dual Image Configuration
- HTTP File Download
- Ping
- TraceRoute
- Ping IPv6
- AutoInstall

2.14.1 Save All Applied Changes

When you click **Submit**, the changes are applied to the system and saved in the running configuration file. However, these changes are not saved to non-volatile memory and will be lost if the system resets. Use the Save All Applied Changes page to make the changes you submit persist across a system reset.

To access the Save All Applied Changes page, click **System > System Utilities > Save All Applied Changes** in the navigation tree.

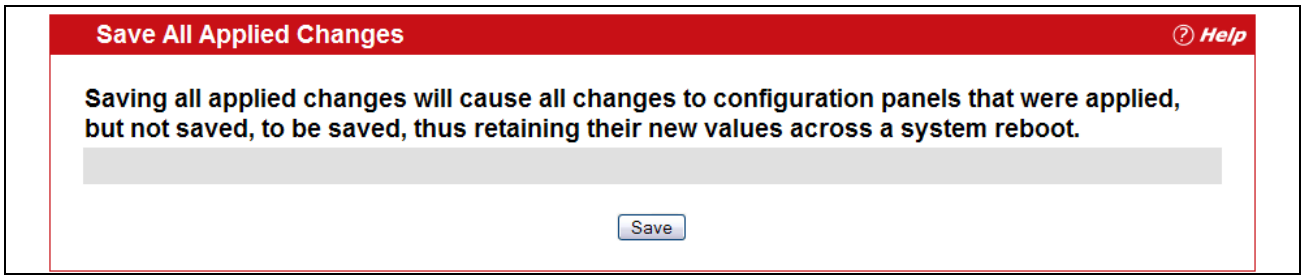

Figure 2-62: Save All Applied Changes

Click **Save** to save all changes applied to the system to NVRAM so that they are retained if the system reboots.

2.14.2 System Reset

Use the System Reset page to reboot the system. If the platform supports stacking, you can reset any of the switches in the stack, or all switches in the stack from this page.

To access the System Reset page, click **System > System Utilities > System Reset** in the navigation tree.

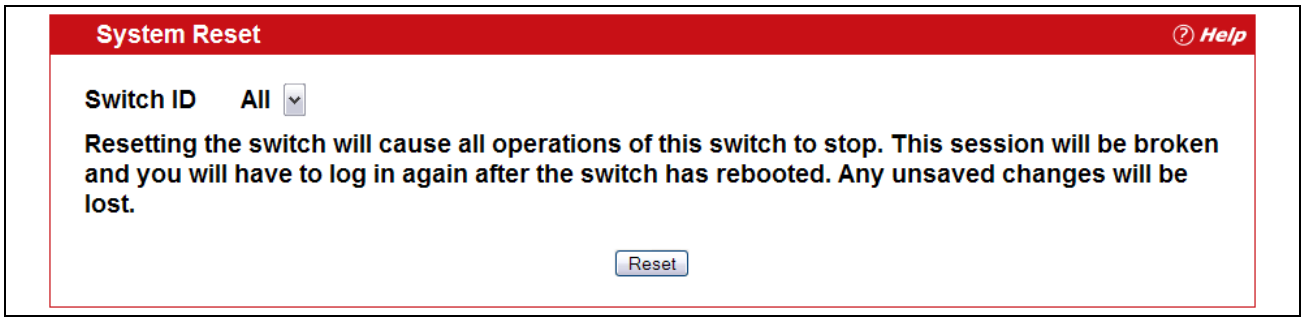

Figure 2-63: System Reset

For Stacking platforms, you can select one or all switches in the stack to reset from the drop-down menu. For platforms that do not support stacking, this field is not present.

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

2.14.3 Reset Configuration to Defaults

Use the Reset Configuration to Defaults page to reset the system configuration to the factory default values.

Note...

By default, the switch does not have an IP address, and the DHCP client is disabled. When you reset the system to its default values, you will not be able to access the Web interface until you connect to the CLI through the serial port and configure network information. For information about configuring network information, see Connecting the Switch to the Network1.

To access the Reset Configuration to Defaults page, click **System > System Utilities > Reset Configuration to Defaults** in the navigation tree.

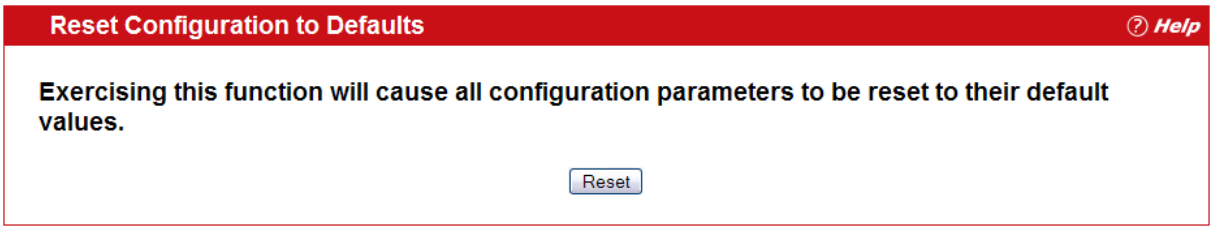

Figure 2-64: Reset Configuration to Defaults

Click **Reset** to restore the factory default settings. The screen refreshes and asks you to confirm the reset. Click **Reset** again to complete the action.

2.14.4 Erase Startup Config File

Use the Erase Startup Config File page to erase the text-based configuration file stored in non-volatile memory. A confirmation screen displays after you select the **Erase** button.

To access the Erase Startup Config File page, click **System > System Utilities > Erase Startup Config File** in the navigation tree.

Figure 2-65: System Reset

Click **Erase** to initiate the process.

2.14.5 Reset Passwords to Defaults

Use the Reset Passwords to Defaults page to reset the passwords for the default read/write (admin) and read-only (guest) users on the system. By default, the passwords are blank. If you have configured additional read-only users on your system, their passwords are not affected.

To access the Reset Passwords to Defaults page, click **System > System Utilities > Reset Passwords to Defaults** in the navigation tree.

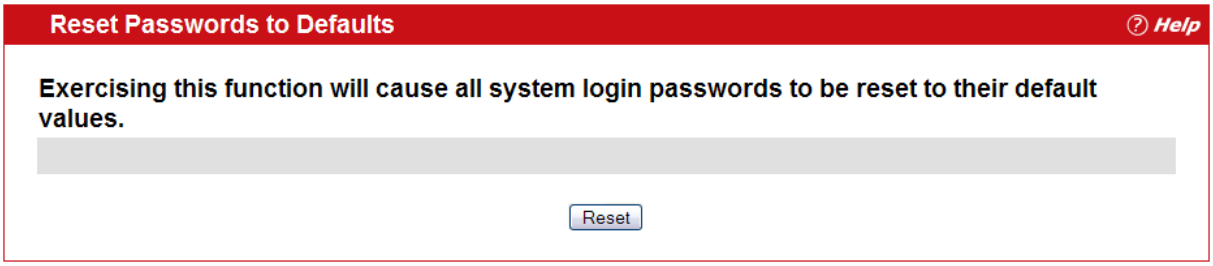

Figure 2-66: Reset Passwords to Defaults

Click **Reset** to restore the passwords for the default users to the factory defaults.

> **Note...**
>
> When the password for the read/write user (admin) changes, you must re-authenticate with the username and default password.

2.14.6 Download File To Switch (TFTP)

Use the Download File to Switch page to download device software, the image file, the configuration files, and SSH or SSL files from a TFTP server to the switch.

You can also download files via HTTP. See 2.14.9HTTP File Download110 for more information.

To access the Download File to Switch page, click **System > System Utilities > Download File to Switch** in the navigation tree.

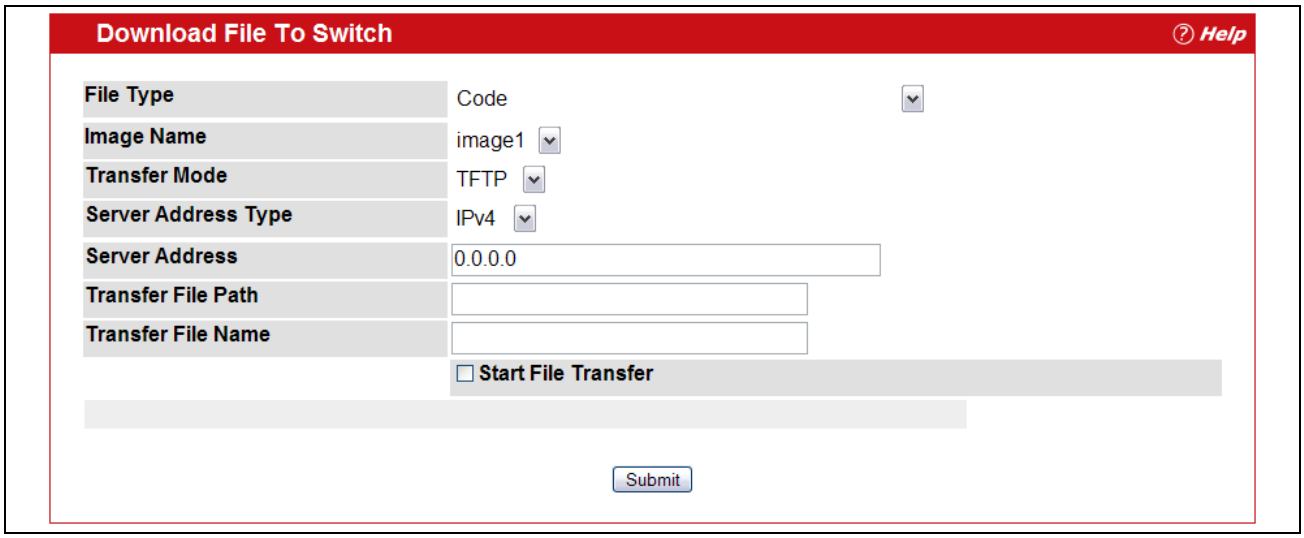

Figure 2-67: Download File to Switch

Table 2-60: Download File to Switch Fields

Field	Description
File Type	Specify what type of file you want to download to the switch: • **CLI Banner**: The CLI banner is the text that displays in the command-line interface before the login prompt. The CLI banner to download is a text file and displays when a user connects to the switch by using telnet, SSH, or a serial connection. • **Code**: The code is the system software image, which is saved in one of two designated files in the file system called images (image1 and image2). The active image stores the active copy; while the other image stores a second copy. The device boots and runs from the active image. If the active image is corrupt, the system automatically boots from the non-active image. This is a safety feature for faults occurring during the boot upgrade process. • **Configuration**: If you have a copy of a valid FASTPATH configuration file (fastpath.cfg) on a TFTP server, you can download it to the switch to overwrite the running and startup configuration files. Upon a successful file transfer, the settings in the configuration file you upload are applied to the switch, and the configuration persists across a system reset. If the file has errors, the update is stopped. The configuration file is not a text file and cannot be edited by using a text editor. • **Text Configuration**: A text-based configuration file enables you to edit a configured text file (startup-config) offline as needed without having to translate the contents for FASTPATH to understand. The most common usage of text-based configuration is to upload a working configuration from a device, edit it offline to personalize it for another similar device (i.e., change the device name, serial number, IP address, etc.), and download it to that device. • **SSH-1 RSA Key File**: SSH-1 Rivest-Shamir-Adleman (RSA) Key File. To download SSH key files, SSH must be administratively disabled and there can be no active SSH sessions. • **SSH-2 RSA Key PEM File:** SSH-2 Rivest-Shamir-Adleman (RSA) Key File (PEM Encoded). To download SSH key files, SSH must be administratively disabled and there can be no active SSH sessions. • **SSH-2 DSA Key PEM File:** SSH-2 Digital Signature Algorithm (DSA) Key File (PEM Encoded). To download SSH key files, SSH must be administratively disabled and there can be no active SSH sessions. • **SSL Trusted Root Certificate PEM File**: SSL Trusted Root Certificate File (PEM Encoded). • **SSL Server Certificate PEM File**: SSL Server Certificate File (PEM Encoded). • **SSL DH Weak Encryption Parameter PEM File**: SSL Diffie-Hellman Weak Encryption Parameter File (PEM Encoded). • **SSL DH Strong Encryption Parameter PEM File**: SSL Diffie-Hellman Strong Encryption Parameter File (PEM Encoded).
Image Name	Specify the code image you want to download, either image1 or image2. This field is only visible when Code is selected as the File Type. The factory default is image1.
Transfer Mode	Specifies the protocol to be used for the transfer: TFTP, SFTP, or SCP.
TFTP Server Address Type	Specify either IPv4,IPv6, or DNS address to indicate the format of the TFTP Server Address field. The factory default is IPv4.
TFTP Server Address	Enter the IP address of the TFTP server in accordance with the format indicated by the TFTP Server Address Type. The factory default is the IPv4 address 0.0.0.0.
TFTP File Path	Enter the path on the TFTP server where the selected file is located. You may enter up to 32 characters. The factory default is blank.
TFTP File Name	Enter the name of the file you want to download from the TFTP server. You may enter up to 32 characters. The factory default is blank.
Start File Transfer	To initiate the download, check this box before clicking **Submit**.

2.14.6.1 Downloading a File to the Switch

Before you download a switch to the file, the following conditions must be true:

- The file to download from the TFTP server is on the server in the appropriate directory.
- The file is in the correct format.
- The switch has a path to the TFTP server.

Use the following procedures to download a file from a TFTP server to the switch.

1. From the **File Type** field, select the type of file to download.
2. If you are downloading a FASTPATH image (Code), select the image on the switch to overwrite. If you are downloading another type of file, the **Image Name** field is not available.

> **Note...**
>
> It is recommended that you not overwrite the active image.

3. Verify the IP address of the TFTP server and ensure that the software image or other file to be downloaded is available on the TFTP server.
4. Complete the **TFTP Server IP Address** and **TFTP File Name** (full path without TFTP server IP address) fields.
5. Click the Start File Transfer check box, and then click **Submit**.

   After you click **Submit**, the screen refreshes and a "File transfer operation started" message appears. After the software is downloaded to the device, a message appears indicating that the file transfer operation completed successfully.

To activate a software image that you download to the switch, see 2.14.8Dual Image Configuration109.

2.14.7 Upload File From Switch (TFTP)

Use the Upload File from Switch page to upload configuration (ASCII) and image (binary) files from the switch to the TFTP server.

To display the Upload File from Switch page, click **System > System Utilities > Upload File from Switch** in the navigation tree.

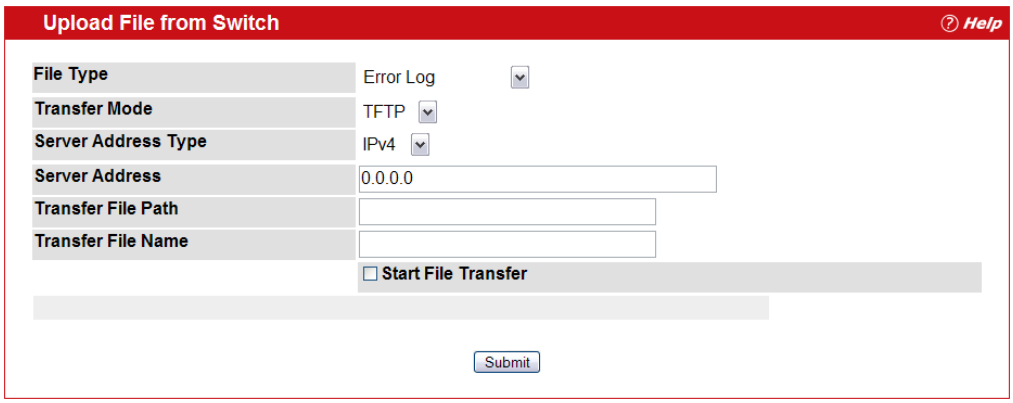

Figure 2-68: Upload File from Switch

Table 2-61: Upload File from Switch Fields

Field	Description
File Type	Specify what type of file you want to upload: • **CLI Banner**: Retrieves the CLI banner file. • **Code**: Retrieves a stored code image. • **Configuration**: Retrieve the stored startup configuration (.cfg) and copy it to a TFTP server. • **Text Configuration**: Retrieves the text configuration file startup-config. • **Error Log**: Retrieves the system error (persistent) log, sometimes referred to as the event log. • **Buffered Log**: Retrieves the system buffered (in-memory) log. • **Trap Log**: Retrieves the system trap records.
Image Name	Specify the code image to upload, either image1 or image2. This field is only visible when Code is selected as the File Type. The factory default is image1.
TFTP Server Address Type	Specify either IPv4 or IPv6 address to indicate the format of the TFTP Server Address field. The factory default is IPv4.
TFTP Server Address	Enter the IP address of the TFTP server in accordance with the format indicated by the TFTP Server Address Type. The factory default is the IPv4 address 0.0.0.0.
TFTP File Path	Enter the path on the TFTP server where you want to put the file. You may enter up to 32 characters. The factory default is blank.
TFTP File Name	Enter a destination file name for the file to upload. You may enter up to 32 characters. The factory default is blank.
Start File Transfer	To initiate the file upload, check this box before clicking **Submit**.

2.14.7.1 Uploading Files

Use the following procedures to upload a file from a TFTP server to the switch.

1. From the **File Type** field, select the type of file to copy from the switch to the TFTP server.
2. If you are uploading a FASTPATH image (Code), select the image on the switch to upload. If you are uploading another type of file, the **Image Name** field is not available.
3. Complete the **TFTP Server Address Type**, **TFTP Server IP Address,** and **TFTP File Name** (full path without TFTP server IP address) fields.
4. Click the **Start File Transfe**r check box, and then click **Submit**.

   After you click **Submit**, the screen refreshes and a “File transfer operation started” message appears. After the software is downloaded to the device, a message appears indicating that the file transfer operation completed successfully.

2.14.8 Dual Image Configuration

The system maintains two versions of the FASTPATH software in permanent storage. One image is the active image, and the second image is the backup image. The active image is loaded during subsequent switch restarts. This feature reduces switch down time when upgrading/downgrading the FASTPATH software.

The system running an older software version will ignore (not load) a configuration file created by the newer software version. When a configuration file created by the newer software version is discovered by the system running an older version of the software, the system will display an appropriate warning to the user.

Use the Dual Image Configuration page to set the boot image.

To display the Dual Image Configuration page, click **System > System Utilities > Dual Image Configuration** in the navigation menu.

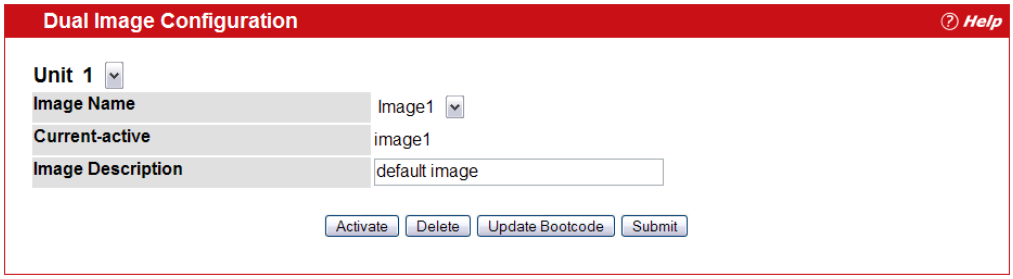

Figure 2-69: Dual Image Configuration

The Active Image page contains the following fields:

Table 2-62: Dual Image Configuration Fields

Field	Description
Image Name	Select image1 or image2 from the drop-down menu to display or configure information about that software image.
Current Active	Displays name of current active image.
Image Description	If desired, enter a descriptive name for the software image.

Click **Activate** to make the image that is selected in the **Image Name** field the next active image for subsequent reboots.

Note...

After activating an image, you must perform a system reset of the switch in order to run the new code.

- Click **Delete** to remove the selected image from permanent storage on the switch.You cannot delete the active image.
- If the file you uploaded contains the boot loader code only, click **Update Bootcode**.
- Click **Submit** to update the image description on the switch.

2.14.9 HTTP File Download

Use the HTTP File Download page to download files of various types to the switch using an HTTP session (i.e., via your web browser).

To display this page, click **System > System Utilities > HTTP File Download** in the navigation menu.

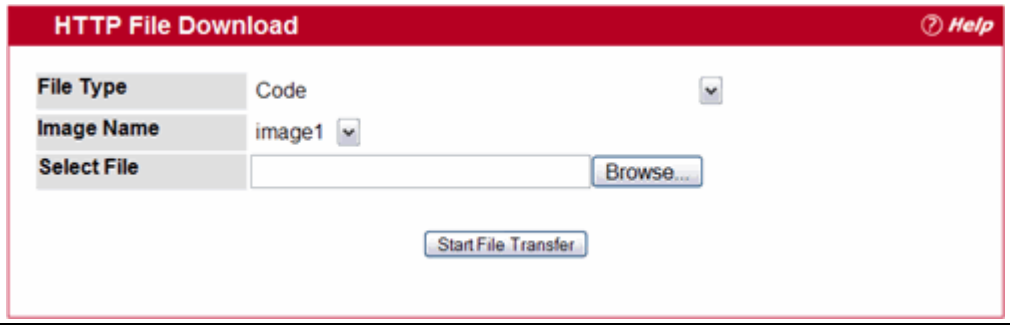

Figure 2-70: HTTP File Download

Table 2-63: HTTP File Download Fields

Field	Description
File Type	Specify the type of file you want to download: • **Code**: Choose this option to upgrade the operational software in flash (default). • **Configuration**: Choose this option to update the switch's configuration. If the file has errors the update will be stopped. • **SSH-1 RSA Key File**: SSH-1 Rivest-Shamir-Adleman (RSA) Key File • **SSH-2 RSA Key PEM File**: SSH-2 Rivest-Shamir-Adleman (RSA) Key File (PEM Encoded) • **SSH-2 DSA Key PEM File**: SSH-2 Digital Signature Algorithm (DSA) Key File (PEM Encoded) • **SSL Trusted Root Certificate PEM File**: SSL Trusted Root Certificate File (PEM Encoded) • **SSL Server Certificate PEM File**: SSL Server Certificate File (PEM Encoded) • **SSL DH Weak Encryption Parameter PEM File**: SSL Diffie-Hellman Weak Encryption Parameter File (PEM Encoded) • **SSL DH Strong Encryption Parameter PEM File**: SSL Diffie-Hellman Strong Encryption Parameter File (PEM Encoded) • **CLI Banner**: Choose this option to download a banner file to be displayed before the login prompt appears. **Note:** To download SSH key files, SSH must be administratively disabled and there can be no active SSH sessions.
Image Name	Specify the code image you want to download, either image1 (the default) or image2. This field is only visible when Code is selected as the File Type.
Select File	Enter the path and filename or browse for the file you want to download. You may enter up to 80 characters.

Click the **Start File Transfer** button to initiate the file download.

2.14.10Ping

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

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

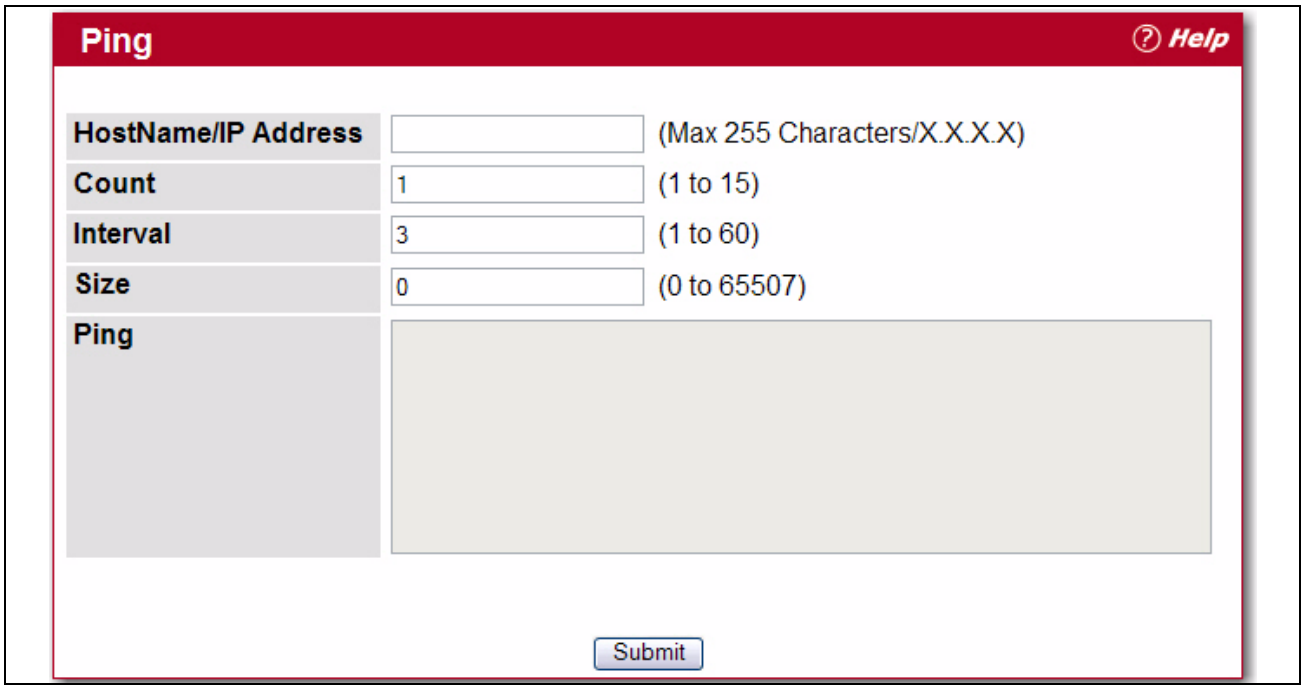

Figure 2-71: Ping

Table 2-64: Ping Fields

Field	Description
Hostname/IP Address	Enter the IP address or the host name of the station you want the switch to ping. The initial value is blank. This information is not retained across a power cycle.
Count	Specify the number of pings to send.
Interval	Specify the number of seconds between pings sent.
Size	Specify the size of the ping packet to send.
Ping	Displays the results of the ping.

Click **Submit** to send the ping. If successful, the results display as shown in Figure 2-72.

2.14.11TraceRoute

You can use the TraceRoute utility to discover the paths that a packet takes to a remote destination.

To display this page, click **System > System Utilities> TraceRoute** in the navigation tree.

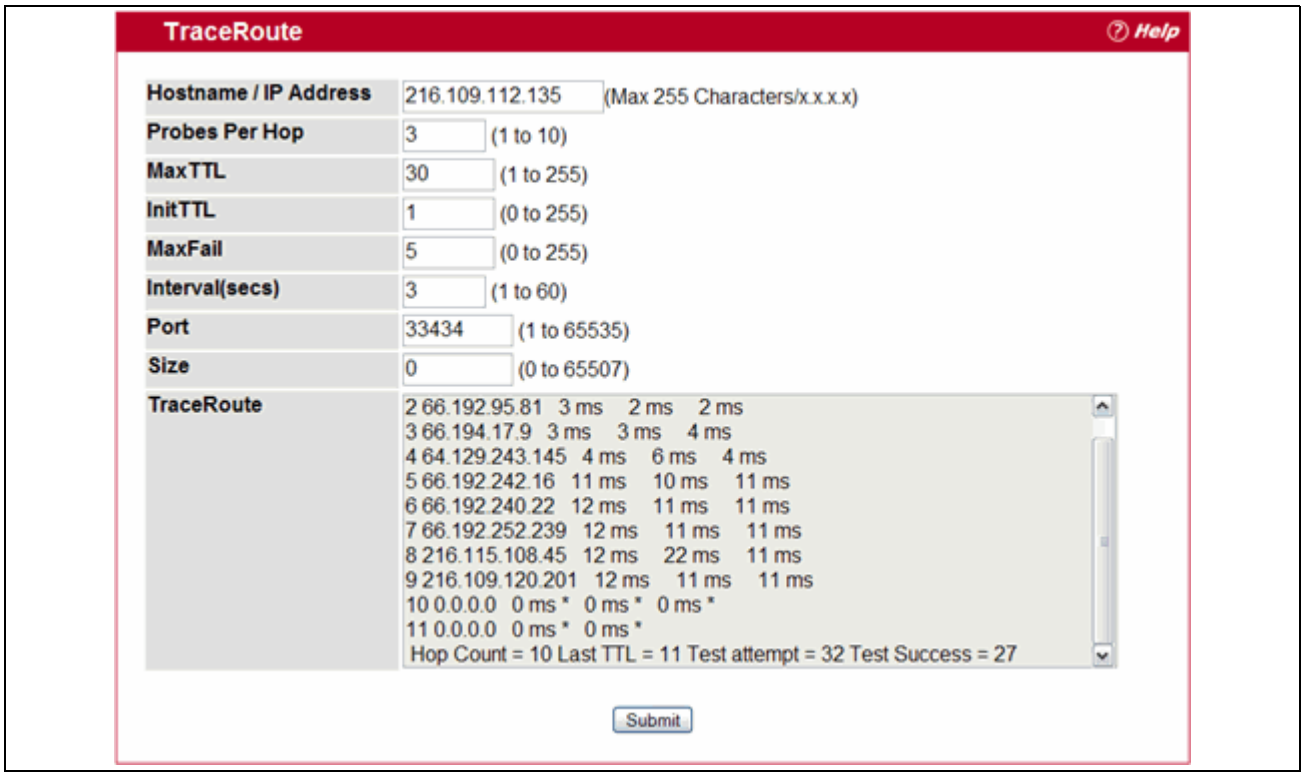

Figure 2-72: TraceRoute

Table 2-65: TraceRoute Fields

	Definition
Hostname/IP Address	Enter the IP address or the hostname of the station you want the switch to discover path for.
Probes Per Hop	Enter the number of times each hop should be probed.
MaxTTL	Enter the maximum time-to-live for a packet in number of hops.
InitTTL	Enter the initial time-to-live for a packet in number of hops.
MaxFail	Enter the maximum number of failures allowed in the session.
Interval	Enter the time between probes in seconds.
Port	Enter the UDP destination port in probe packets.
Size	Enter the size of probe packets.
TraceRoute	Displays the output from a traceroute.

Click **Submit** to initiate the traceroute. The results display in the TraceRoute box.

2.14.12Ping IPv6

Use the Ping IPv6 page to send a ping request to a specified IPv6 address. You can use this feature to check whether the switch can communicate with a particular IPv6 station.

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

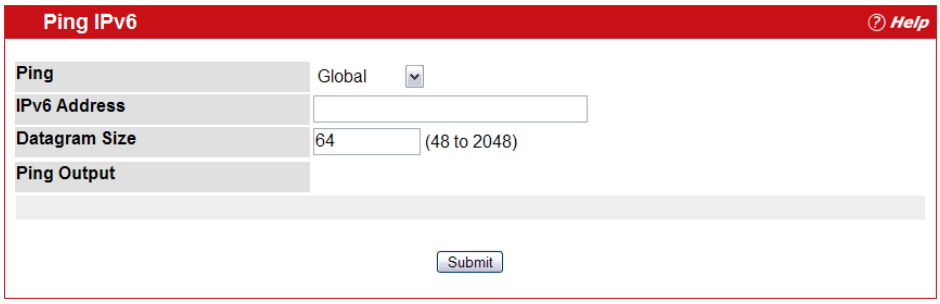

Figure 2-73: Ping IPv6

Table 2-66: Ping IPv6 Fields

Field	Description
Ping	Select a the type of address to ping. • **Global**: Ping a global IPv6 Address. A global IPv6 address can be routed beyond the local network and can be reached on the IPv6 Internet. • **Link Local**: Ping a link local IPv6 address. A link local address always begins with FE80::/64. IPv6 routers do not forward link-local traffic beyond the link.
Interface	This field displays only when **Link Local** is selected. Select an IPv6 interface to initiate the ping.
IPv6 Address	This field displays only when **Global** is selected. Enter the IPv6 address of the station you want the switch to ping. The initial value is blank. The IPv6 Address you enter is not retained across a power cycle.
Link Local Address	This field is present if you select **Link Local** from the **Ping** field. Enter the link local address of the station you want the switch to ping. The initial value is blank. The Link Local Address you enter is not retained across a power cycle.
Datagram Size	Enter the datagram size. The valid range is platform-dependent.
Ping Output	The results display in this field.

Click **Submit** to send the specified number of pings. The results display in the Ping Output box.

2.14.13AutoInstall

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

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

- The IP address and subnet mask (option 1) to be assigned to the switch.
- The IP address of a default gateway (option 3), if needed for IP communication.
- The identification of the TFTP server from which to obtain the boot file. This is given by any of the following fields, in the priority shown (highest to lowest):
  - The sname field of the DHCP reply.
  - The hostname of the TFTP server (option 66). Either the TFTP address or name is specified—not both—in most network configurations. If a TFTP hostname is given, a DNS server is required to translate the name to an IP address.

- The IP address of the TFTP server (option 150).
- The address of the TFTP server supplied in the siaddr field.
- The name of the configuration file (boot file or option 67) to be downloaded from the TFTP server. **The boot file name must have a file type of *.cfg.**

- The IP addresses of DNS name servers (option 6). The IP addresses of DNS name servers should be returned from the DHCP server only if the DNS server is in the same LAN as the switch performing AutoInstall. A DNS server is needed to resolve the IP address of the TFTP server if only the "sname" or option 66 values are returned to the switch.

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

To display the AutoInstall page, click **System > System Utilities> AutoInstall**.

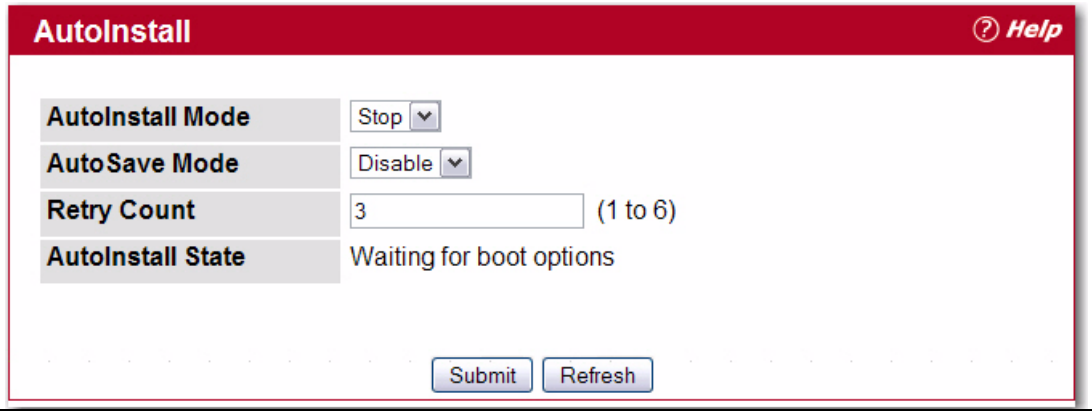

Figure 2-74: AutoInstall

Table 2-67: AutoInstall

Field	Definition
AutoInstall Mode	Select **Start** to initiate sending a request to a DHCP server to obtain an IP address of a server and the configuration file name. If it obtains the server address, AutoInstall proceeds to search for and download a configuration file from the server. If successful, it applies the configuration file to the switch. After starting the AutoInstall process, you can monitor the status of the process by the messages in the AutoInstall State and Retry Count fields. After 3 retries, AutoInstall informs the failure to the TR-069 module, and TR-069 client tries to download the configuration file from a TFTP server through RequestDownload RPC. You can click **Stop** to end the process.
AutoSave Mode	Enable or disable saving the network configuration to non-volatile memory. When enabled, the configuration is saved after downloading from the TFTP server without operator intervention. When disabled, the operator must explicitly save the configuration, if needed.
Retry Count	The number of times the switch has attempted to contact the TFTP server during the current AutoInstall session.
AutoInstall State	The status of the current or most recently completed AutoInstall session.

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

2.15 Managing SNMP Traps

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

2.15.1 Trap Flags

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

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

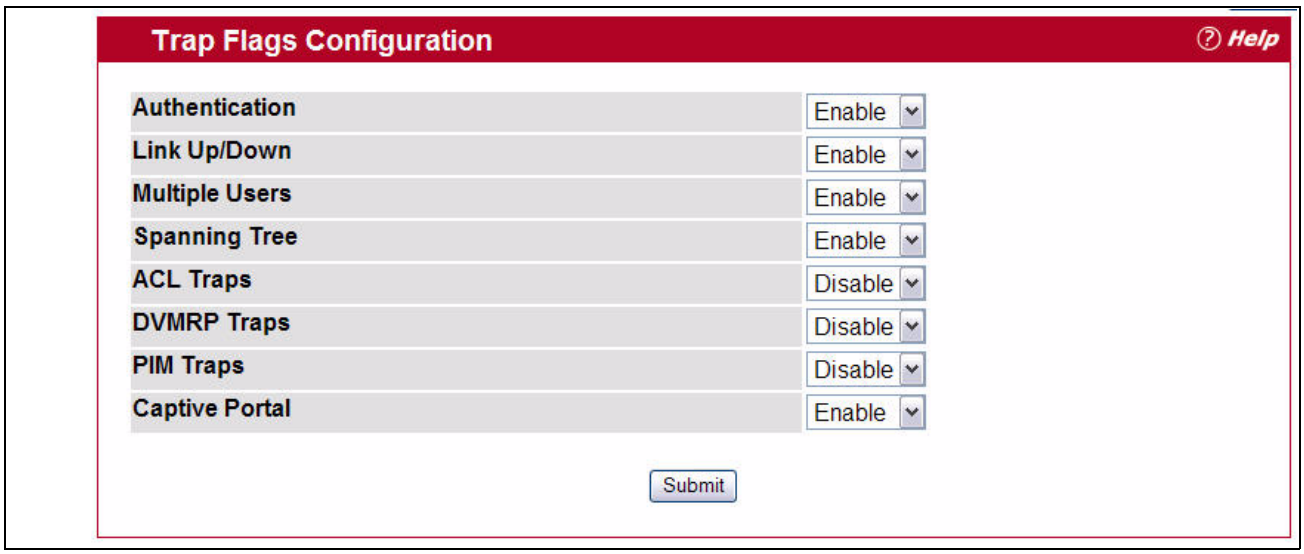

Figure 2-75: Trap Flags Configuration

The fields available on the Trap Flags page depends on the packages installed on your system. For example, if your system does not have the BGP4 package installed, the BGP Traps field is not available. Figure 2-75 and Table 2-68show the fields that are available on a system with all packages installed.

Table 2-68: Trap Flags Configuration Fields

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

Table 2-68: Trap Flags Configuration Fields

Field	Description
BGP4 Traps	Enable or disable activation of BGP traps by selecting the corresponding line on the pulldown entry field. The factory default is enabled. If your system does not support BGP, this field is not available.
DVMRP Traps	Enable or disable activation of DVMRP traps by selecting the corresponding line on the pulldown entry field. The factory default is enabled. If your system does not support Multicast, this field is not available.
OSPF Traps	Enable or disable activation of OSPF traps by selecting the corresponding line on the pulldown entry field. The factory default is enabled. This field can be configured only if the OSPF admin mode is enabled. If your system does not support OSPF routing, this field is not available.
PIM Traps	Enable or disable activation of PIM traps by selecting the corresponding line on the pulldown entry field. The factory default is enabled. If your system does not support Multicast, this field is not available.
Captive Portal	Select Enable to allow the SNMP agent on the switch to generate captive portal SNMP traps that are enabled. Select Disable to prevent the SNMP agent on the switch from generating any captive portal SNMP traps, even if they are individually enabled.

If you make any changes to this page, click **Submit** to apply the changes to the system.

2.15.2 OSPFv2 Trap Flags

Use the OSPFv2 Trap Flags page to enable or disable OSPFv2 traps the switch can send to an SNMP manager. When the condition identified by an active trap is encountered by the switch, a trap message is sent to any enabled SNMP Trap Receivers, and a message is written to the trap log.

To access the OSPFv2 Trap Flags page, click **System > Trap Manager > OSPFv2 Trap Flags** page.

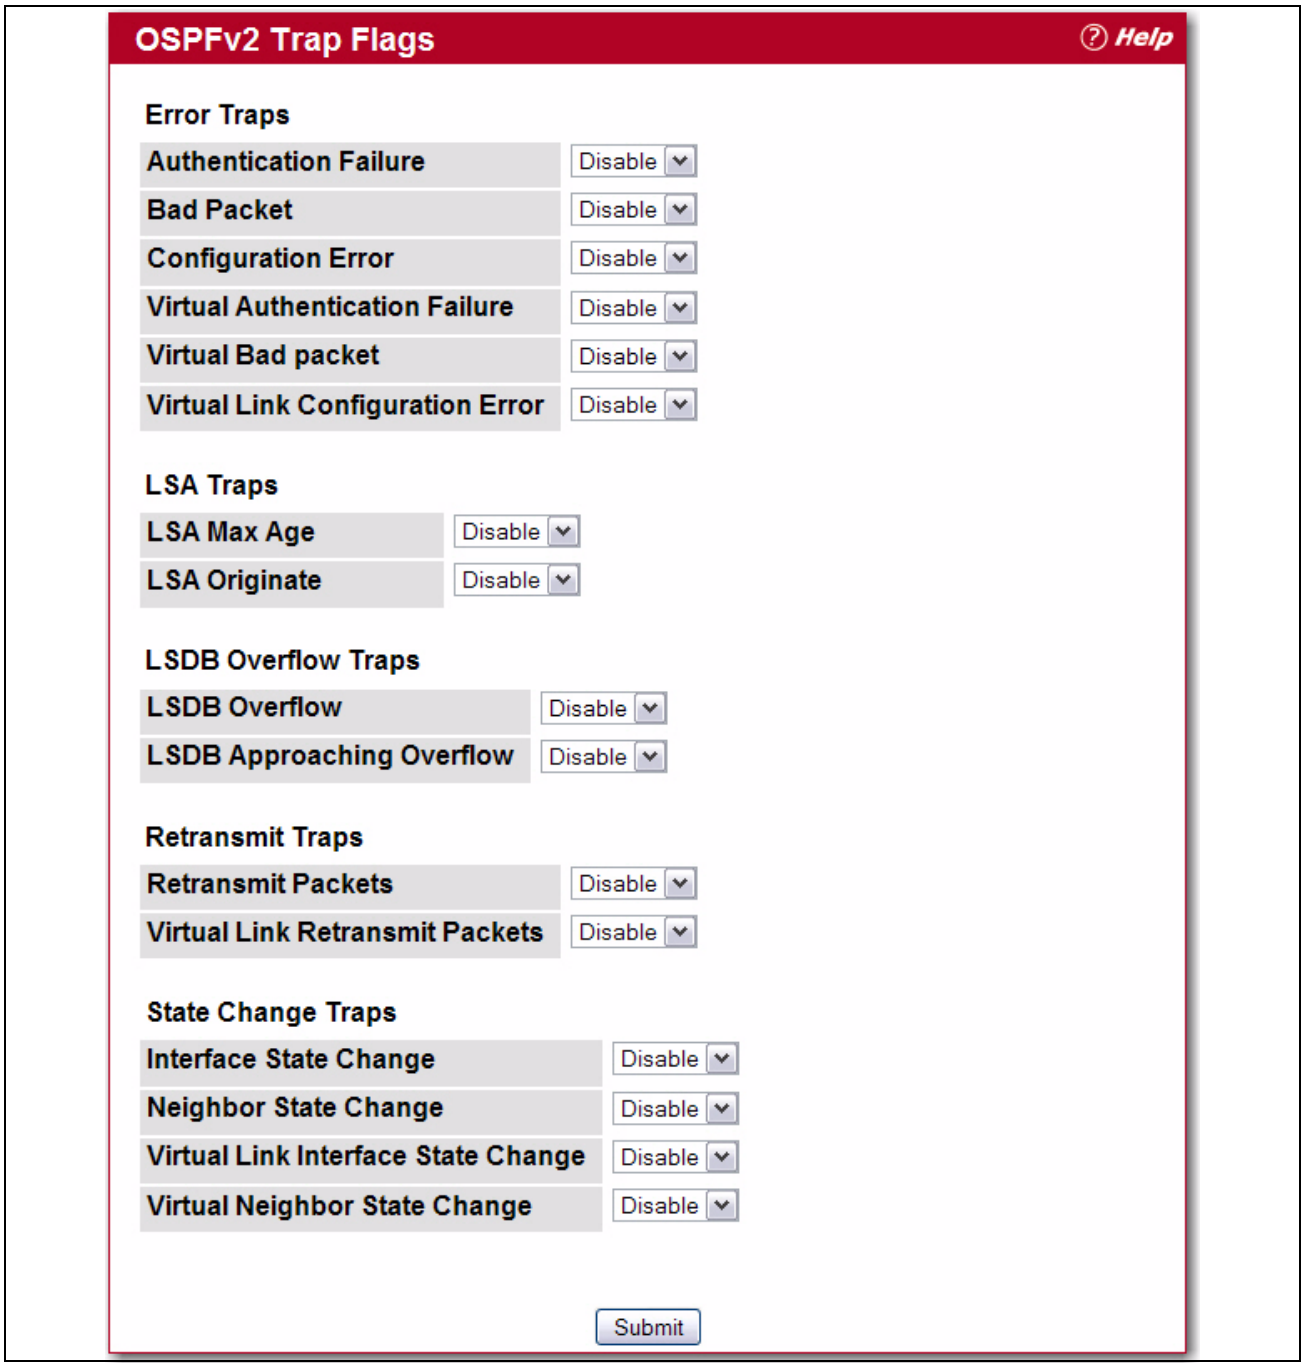

Figure 2-76: OSPFv2 Trap Flags Configuration

Table 2-69: OSPFv2 Trap Flags Configuration Fields

Field	Description
Error Traps	
Authentication Failure	Signifies that a packet has been received on a non-virtual interface from a router whose authentication key or authentication type conflicts with this router's authentication key or authentication type. The factory default is disabled.
Bad Packet	Signifies that an OSPF packet has been received on a non-virtual interface that cannot be parsed. The factory default is disabled.
Configuration Error	Signifies that a packet has been received on a non-virtual interface from a router whose configuration parameters conflict with this router's configuration parameters. The factory default is disabled.
Virtual Authentication Failure	Signifies that a packet has been received on a virtual interface from a router whose authentication key or authentication type conflicts with this router's authentication key or authentication type. The factory default is disabled.
Virtual Bad packet	Signifies that an OSPF packet has been received on a virtual interface that cannot be parsed. The factory default is disabled.
Virtual Link Configuration Error	Signifies that a packet has been received on a virtual interface from a router whose configuration parameters conflict with this router's configuration parameters. The factory default is disabled.
LSA Traps	
LSA Max Age	Signifies that one of the LSA in the router's link-state database has aged to MaxAge. The factory default is disabled.
LSA Originate	Signifies that a new LSA has been originated by this router. This trap should not be invoked for simple refreshes of LSAs (which happens every 30 minutes), but instead will only be invoked when an LSA is (re)originated due to a topology change. Additionally, this trap does not include LSAs that are being flushed because they have reached MaxAge. The factory default is disabled.
LSDB Overflow Traps	
LSDB Overflow	Signifies that the number of LSAs in the router's link-state database has exceeded OSPF External LSDB Limit. The factory default is disabled.
LSDB Approaching Overflow	Signifies that the number of LSAs in the router's link-state database has exceeded ninety percent of OSPF External LSDB Limit. The factory default is disabled.
Retransmit Traps	
Retransmit Packets	Signifies that an OSPF packet has been retransmitted on a non-virtual interface. All packets that may be retransmitted are associated with an LSDB entry. The LS type, LS ID, and Router ID are used to identify the LSDB entry. The factory default is disabled.
Virtual Link Retransmit Packets	Signifies that an OSPF packet has been retransmitted on a virtual interface. All packets that may be retransmitted are associated with an LSDB entry. The LS type, LS ID, and Router ID are used to identify the LSDB entry. The factory default is disabled.
State Change Traps	

Table 2-69: OSPFv2 Trap Flags Configuration Fields

Field	Description
Interface State Change	Signifies that there has been a change in the state of a non-virtual OSPF interface. This trap should be generated when the interface state regresses (e.g., goes from Dr to Down) or progresses to a terminal state (i.e., Point-to-Point, DR Other, Dr, or Backup). The factory default is disabled.
Neighbor State Change	Signifies that there has been a change in the state of a non-virtual OSPF neighbor. This trap should be generated when the neighbor state regresses (e.g., goes from Attempt or Full to 1-Way or Down) or progresses to a terminal state (e.g., 2-Way or Full). When an neighbor transitions from or to Full on non-broadcast multi-access and broadcast networks, the trap should be generated by the designated router. A designated router transitioning to Down will be noted by OSPF Interface State Change. The factory default is disabled.
Virtual Link Interface State Change	Signifies that there has been a change in the state of an OSPF virtual interface. This trap should be generated when the interface state regresses (e.g., goes from Point-to-Point to Down) or progresses to a terminal state (i.e., Point-to-Point). The factory default is disabled.
Virtual Neighbor State Change	Signifies that there has been a change in the state of an OSPF virtual neighbor. This trap should be generated when the neighbor state regresses (e.g., goes from Attempt or Full to 1-Way or Down) or progresses to a terminal state (e.g., Full). The factory default is disabled.

If you make any changes to this page, click **Submit** to apply the changes to the system.

2.15.3 OSPFv3 Trap Flags

Use the OSPFv3 Trap Flags page to enable or disable traps the switch can send to an SNMP manager. When the condition identified by an active trap is encountered by the switch, a trap message is sent to any enabled SNMP Trap Receivers, and a message is written to the trap log.

To access the Trap Flags page, click **System > Trap Manager > OSPFv3 Trap Flags** page.

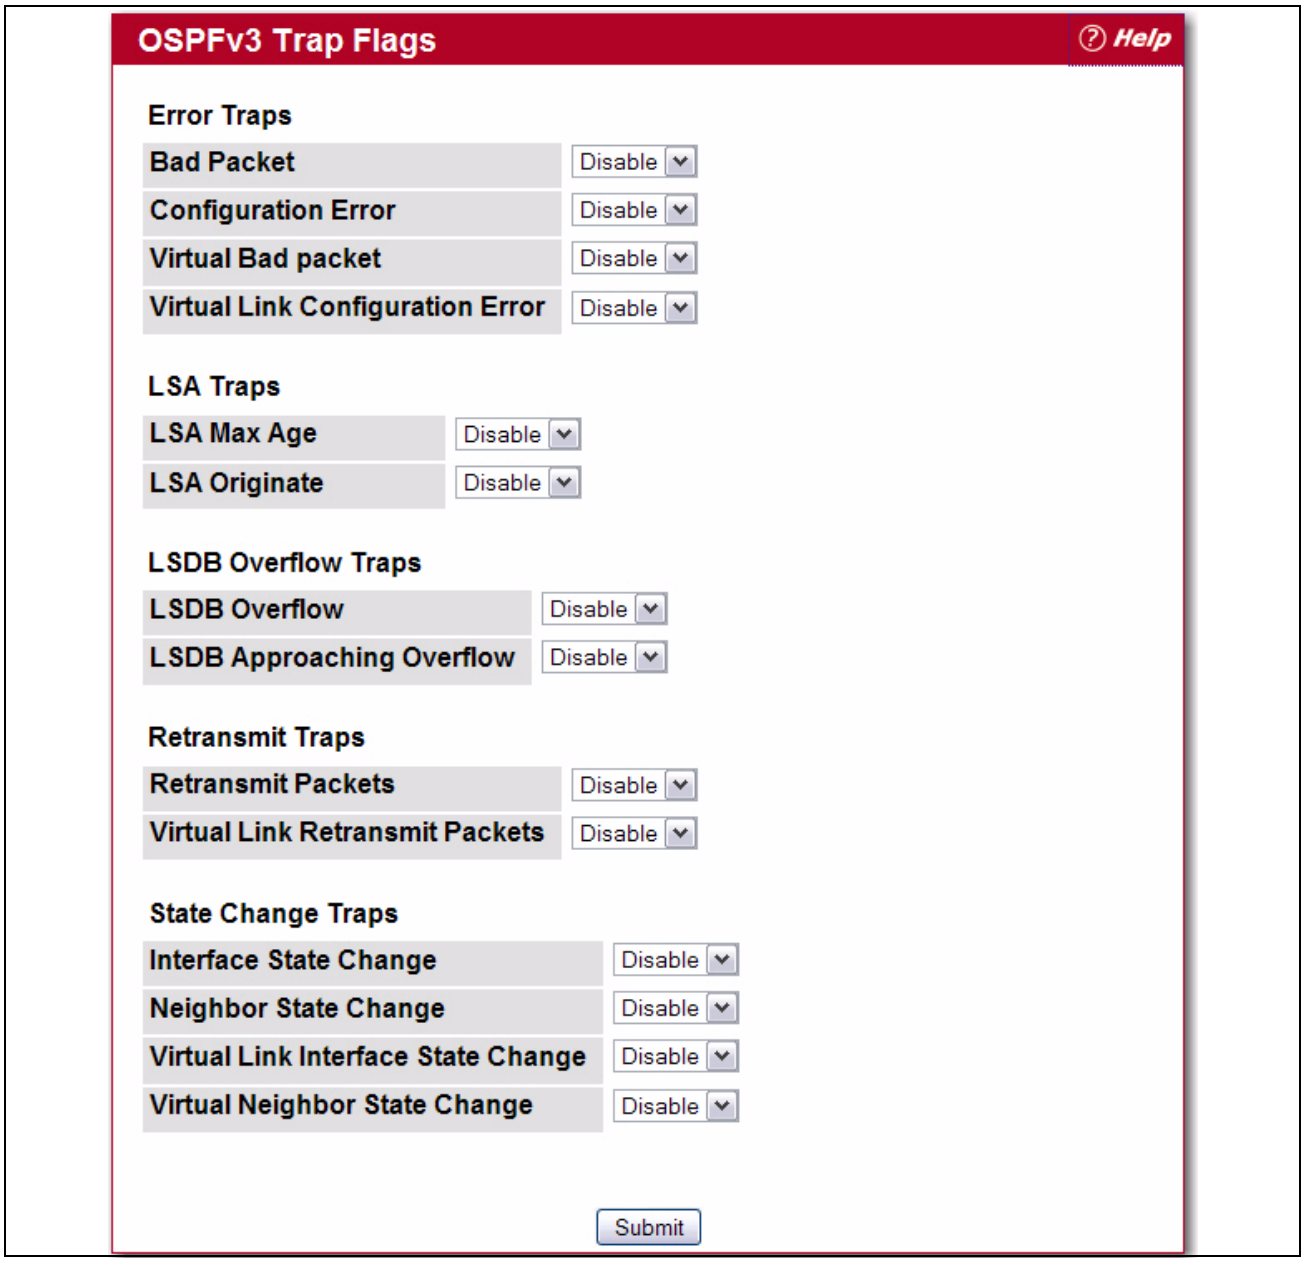

Figure 2-77: Trap Flags Configuration

Table 2-70: OSPFv3 Trap Flags Configuration Fields

Field	Description
Error Traps	
Bad Packet	Signifies that an OSPF packet has been received on a non-virtual interface that cannot be parsed. The factory default is disabled.
Configuration Error	Signifies that a packet has been received on a non-virtual interface from a router whose configuration parameters conflict with this router's configuration parameters. The factory default is disabled.
Virtual Bad packet	Signifies that an OSPF packet has been received on a virtual interface that cannot be parsed. The factory default is disabled.
Virtual Link Configuration Error	Signifies that a packet has been received on a virtual interface from a router whose configuration parameters conflict with this router's configuration parameters. The factory default is disabled.
LSA Traps	
LSA Max Age	Signifies that one of the LSA in the router's link-state database has aged to MaxAge. The factory default is disabled.
LSA Originate	Signifies that a new LSA has been originated by this router. This trap should not be invoked for simple refreshes of LSAs(which happens every 30 minutes), but instead will only be invoked when an LSA is (re)originated due to a topology change. Additionally, this trap does not include LSAs that are being flushed because they have reached MaxAge. The factory default is disabled.
LSDB Overflow Traps	
LSDB Overflow	Signifies that the number of LSAs in the router's link-state database has exceeded OSPF External LSDB Limit. The factory default is disabled.
LSDB Approaching Overflow	Signifies that the number of LSAs in the router's link-state database has exceeded ninety percent of OSPF External LSDB Limit. The factory default is disabled.
Retransmit Traps	
Retransmit Packets	Signifies that an OSPF packet has been retransmitted on a non-virtual interface. All packets that may be retransmitted are associated with an LSDB entry. The LS type, LS ID, and Router ID are used to identify the LSDB entry. The factory default is disabled.
Virtual Link Retransmit Packets	Signifies that an OSPF packet has been retransmitted on a virtual interface. All packets that may be retransmitted are associated with an LSDB entry. The LS type, LS ID, and Router ID are used to identify the LSDB entry. The default is disabled.
State Change Traps	

Table 2-70: OSPFv3 Trap Flags Configuration Fields

Field	Description
Interface State Change	Signifies that there has been a change in the state of a non-virtual OSPF interface. This trap should be generated when the interface state regresses (e.g., goes from Dr to Down) or progresses to a terminal state (i.e., Point-to-Point, DR Other, Dr, or Backup). The factory default is disabled.
Neighbor State Change	Signifies that there has been a change in the state of a non-virtual OSPF neighbor. This trap should be generated when the neighbor state regresses (e.g., goes from Attempt or Full to 1-Way or Down) or progresses to a terminal state (e.g., 2-Way or Full). When an neighbor transitions from or to Full on non-broadcast multi-access and broadcast networks, the trap should be generated by the designated router. A designated router transitioning to Down will be noted by OSPF Interface State Change. The factory default is disabled.
Virtual Link Interface State Change	Signifies that there has been a change in the state of an OSPF virtual interface. This trap should be generated when the interface state regresses (e.g., goes from Point-to-Point to Down) or progresses to a terminal state (i.e., Point-to-Point). The factory default is disabled.
Virtual Neighbor State Change	Signifies that there has been a change in the state of an OSPF virtual neighbor. This trap should be generated when the neighbor state regresses (e.g., goes from Attempt or Full to 1-Way or Down) or progresses to a terminal state (e.g., Full). The factory default is disabled.

If you make any changes to this page, click **Submit** to apply the changes to the system.

2.15.4 Trap Log

Use the Trap Log page to view the entries in the trap log. For information about how to copy the file to a TFTP server, see 2.14.7Upload File From Switch (TFTP)108.

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

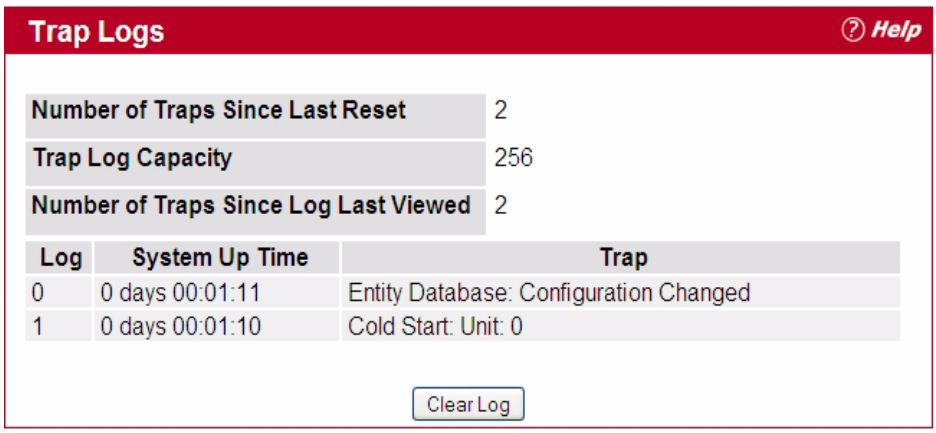

Figure 2-78: Trap Log

Table 2-71: Trap Log Fields

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

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

2.16 Managing the DHCP Server

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

- Global Configuration
- Pool Configuration
- Pool Options
- Reset Configuration
- Bindings Information
- Server Statistics
- Conflicts Information

2.16.1 Global Configuration

Use the Global Configuration page to configure DHCP global parameters.

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

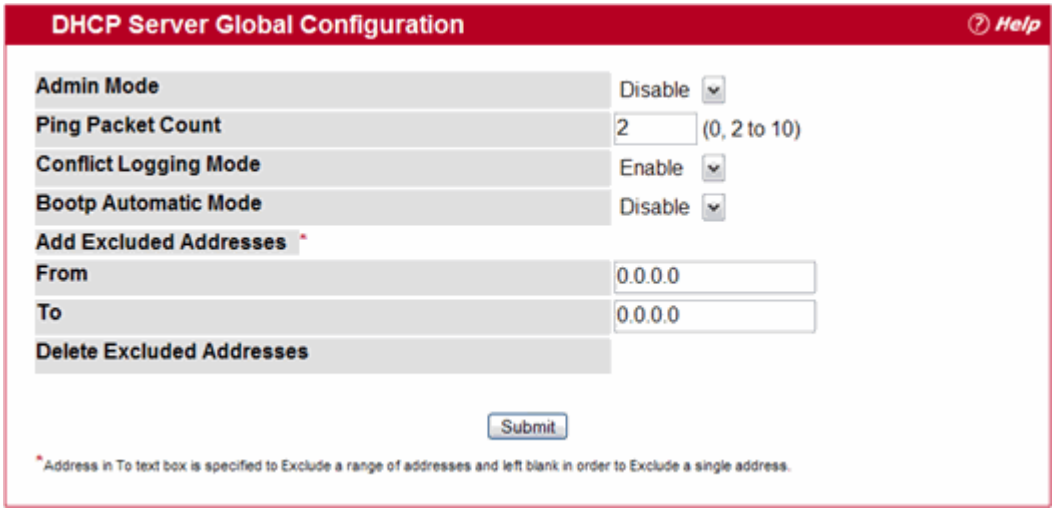

Figure 2-79: DHCP Server Global Configuration

Table 2-72: DHCP Server Global Configuration Fields

Field	Description
Admin Mode	Enables or disables DHCP server operation on the switch. The default value is Disable.
Ping Packet Count	Specifies the number of packets a server sends to a Pool address to check for duplication as part of a ping operation. Default value is 2. The valid range is (0, 2 to 10). Setting the value to 0 disables the function.
Conflict Logging Mode	Specifies whether to enable or disable conflict logging on a DHCP Server. The default value is Enable.
Bootp Automatic Mode	Specifies whether to enable or disable Bootp for dynamic pools.
Enable	Allows the allocation of the addresses in the automatic address pool to the BootP client.
Disable	Does not use the automatic address pool addresses for BootP clients. This is the default value.
Add Excluded Addresses	Use the **From** and **To** fields to specify the IP addresses that the server should not assign to the client. If you want to exclude a range of addresses, set the range boundaries.
From	To exclude an address range, specify the low address in the range.To specify a single address to exclude, enter the address in the **From** field and leave the **To** field at the default value of 0.0.0.0. For example, in Figure 2-80, the user is adding the address 192.168.17.100 to the excluded addresses list.
To	To exclude an address range, specify the high address in the range. To exclude a single address, do not enter a value in this field.
Delete Excluded Addresses	After you add excluded addresses, they appear below this field title, as Figure 2-80 shows. Each address or address range has a check box next to it.

- If you change any settings or add an excluded address range, click **Submit** to apply the changes to the system. Each time you enter a value in the **From** or **To** fields, click **Submit** to add the address or address range to the excluded address list.
- To Delete an address or address range from the excluded address list, select one or more check box beneath the Delete **Excluded Addresses** field and click **Submit**.

2.16.2 Pool Configuration

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

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

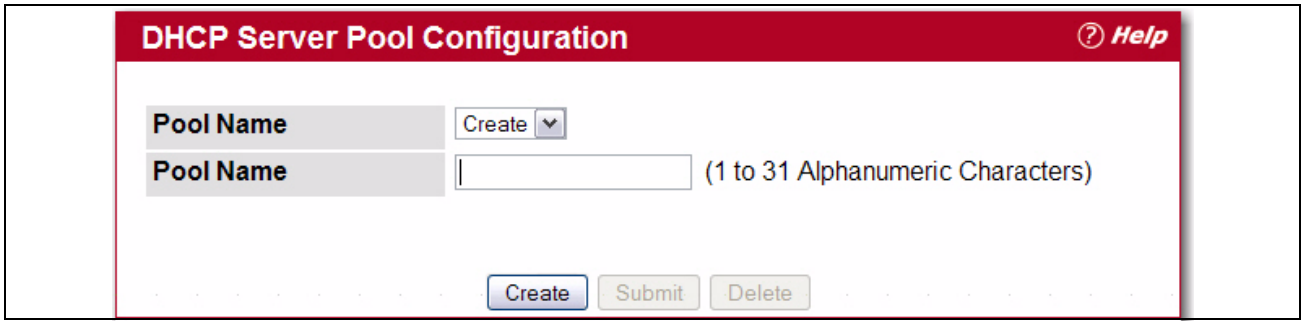

Figure 2-80: Pool Configuration

Table 2-73: Pool Configuration Fields

Field	Description
Pool Name	For a user with read/write permission, this field would show names of all the existing pools along with an additional option Create. When the user selects Create, another text box, Pool Name, appears where the user may enter name for the Pool to be created.For a user with read-only permission, this field would show names of the existing pools only.
Pool Name	This field appears when the user with read-write permission has selected Create in the Drop Down list against Pool Name. Specifies the Name of the Pool to be created. Pool Name can be up to 31 characters in length.

In Figure 2-80, some of the blank fields where you add IP addresses have been edited out of the image for display purposes. You can add up to eight addresses in the Default Router Addresses, DNS Server Addresses, NetBIOS name Server Addresses and IP Address Value fields.

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

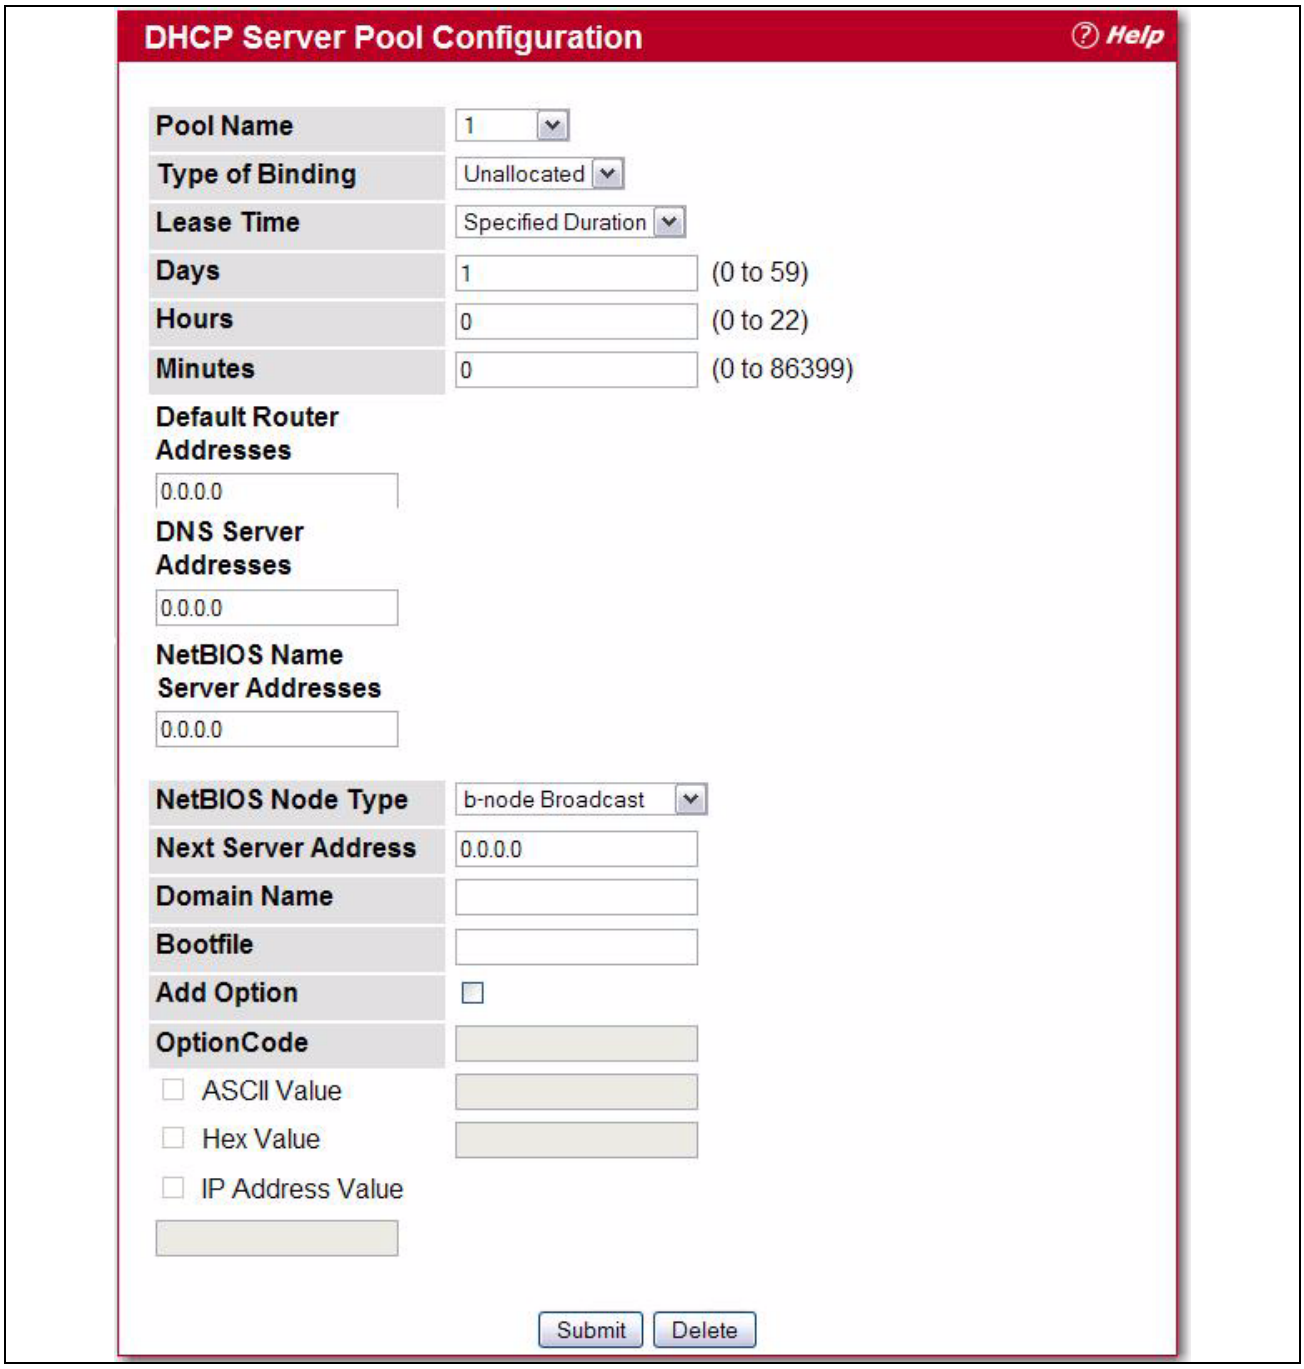

Figure 2-81: Pool Configuration Continued

Table 2-74: Pool Configuration Fields

Field	Description
Pool Name	For a user with read/write permission, this field would show names of all the existing pools along with an additional option Create. When the user selects Create, another text box, Pool Name, appears where the user may enter name for the Pool to be created.For a user with read-only permission, this field would show names of the existing pools only.
Pool Name	This field appears when the user with read-write permission has selected Create in the Drop Down list against Pool Name. Specifies the Name of the Pool to be created. Pool Name can be up to 31 characters in length.
Type of Binding	Specifies the type of binding for the pool. • **Unallocated**: The addresses are not assigned to a client. • **Automatic**: The IP address is automatically assigned to a client by the DHCP server. • **Manual**: You statically assign an IP address to a client based on the client's MAC address.
Network Number	If you specify Dynamic as the type of binding, this field appears. Specifies the network number (host bits) for a DHCP address of a dynamic pool. For example, if 192.168.5.0 is the network number and 255.255.255.0 is the network mask (or a prefix length of 24) for the pool, the IP addresses in the pool range from 192.168.5.1 - 192.168.5.254.
Network Mask	For dynamic bindings, this field specifies the subnet mask for a DHCP address of a dynamic pool. You can enter a value in Network Mask **or** Prefix Length to specify the subnet mask, but do not enter a value in both fields.
Prefix Length	For dynamic bindings, this field specifies the subnet number for a DHCP address of a dynamic pool. You can enter a value in Network Mask **or** Prefix Length to specify the subnet mask, but do not enter a value in both fields. The valid range is 0 to 32.
Client Name	For manual bindings, this field specifies a name for the client to which the DHCP server will statically assign an IP address. This field is optional.
Hardware Address	For manual bindings, this field specifies the MAC address of the hardware platform of the DHCP client.
Hardware Address Type	For manual bindings, this field specifies the protocol of the hardware platform of the DHCP client. Valid types are ethernet and ieee802. Default value is ethernet.
Client ID	For manual bindings, this field specifies the Client Identifier for DHCP manual Pool.
Host Number	For manual bindings, this field specifies the IP address to be statically assigned to a DHCP client. The host can be set only if at least one among of Client Identifier or Hardware Address is specified. Deleting Host would delete Client Name, Client ID, Hardware Address for the Manual Pool and set the Pool Type to Unallocated.
Host Mask	For manual bindings, this field specifies the subnet mask to be statically assigned to a DHCP client. You can enter a value in Host Mask **or** Prefix Length to specify the subnet mask, but do not enter a value in both fields.
Prefix Length	For manual and dynamic bindings, this field specifies the subnet mask for a manual binding to a DHCP client. You can enter a value in Network Mask **or** Prefix Length to specify the subnet mask, but do not enter a value in both fields. The valid range is 0 to 32.
Lease Time	Specifies the type of lease to assign clients: • **Infinite**: For dynamic bindings, an infinite least time is a lease period of 60 days. For manual bindings, an infinite lease time means the lease period does not expire. • **Specified Duration**: Allows you to specify the lease period. The default value is Specified Duration. • **Db-node Broadcast**: Uses broadcasted queries.
Days	For a Specified Duration lease time, this field specifies the number of days for the lease period. The default value is 1, and the valid range is 0-59.

Table 2-74: Pool Configuration Fields (Continued)

Field	Description
Hours	For a Specified Duration lease time, this field specifies the number of hours for the lease period. The default value is 1, and the valid range is 0-1439.
Minutes	For a Specified Duration lease time, this field specifies the number of minutes for the lease period. The default value is 1, and the valid range is 0-86399.
Default Router Addresses	Specifies the list of default router IP addresses for the pool. You can specify up to eight addresses in order of preference.
DNS Server Addresses	Specifies the list of DNS server IP addresses for the pool. You can specify up to eight addresses in order of preference.
NetBIOS Name Server Addresses	Specifies the list of NetBIOS name server IP addresses for the pool. You can specify up to eight addresses in order of preference.
NetBIOS Node Type	Specifies the NetBIOS node type for DHCP clients: • **p-node Peer-to-Pee**r: Uses point-to-point name queries to a name server. • **m-node Mixed**: Uses broadcasts first, then uses queries the name server. • **h-node Hybrid**: Uses queries the name server first, and then uses broadcasts.
Next Server Address	Specifies the IP address of the next server in the client's boot process, such as a TFTP server.
Domain Name	Specifies the domain name for a DHCP client. The domain name can be up to 255 characters in length.
Bootfile	Specifies the name of the default boot image for a DHCP client. The file name can be up to 128 characters in length.
Add Options	The rest of the fields on the page allow you to add and configure DHCP options. See RFC 2132 for more information about DHCP options.
Code	Specifies the DHCP option code. The valid range is 1-254.
Ascii Value	Specifies an NVT ASCII character string.
Hex Value	Specifies dotted hexadecimal data. Each byte in hexadecimal character strings is 2 hexadecimal digits. Each byte can be separated by a colon or white space. A period separates 2 bytes/4 hexadecimal digits.
IP Address Values	Specifies the Option IP addresses.

- After you configure values for the DHCP address pool, click **Submit** to create the pool and apply the changes to the system.
- To delete a pool, select the pool from the **Pool Name** drop-down menu and click **Delete**.

2.16.3 Pool Options

Use the Pool Options page to configure DHCP options that the DHCP server can pass to the client. For more information about DHCP options, see RFC 2132.

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

If no DHCP pools exist, the Pool Options page does not display the fields shown in Figure 2-82.

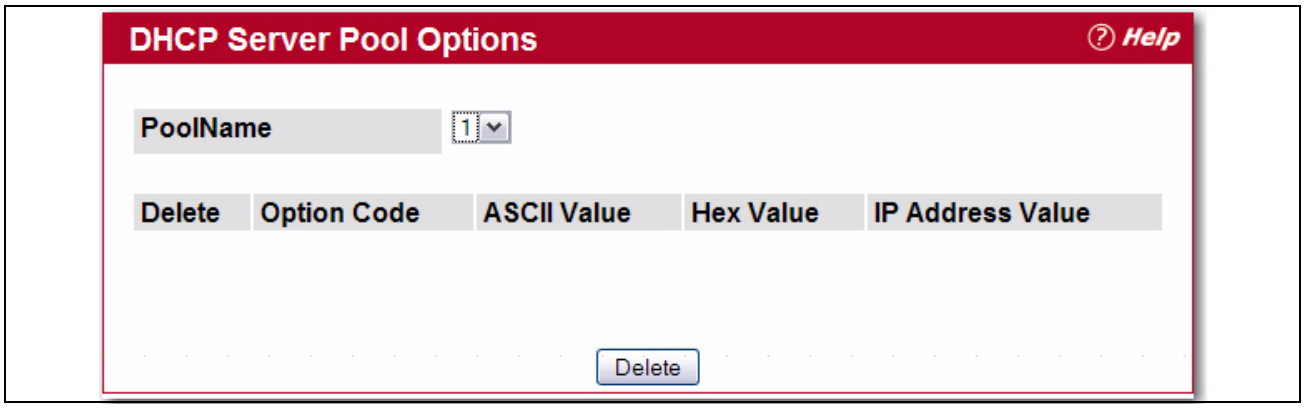

Figure 2-82: Pool Options

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

Table 2-75: Pool Options Fields

Field	Description
Pool Name	Select the DHCP pool to with the options you want to view or configure.
Option Code	Displays the DHCP option code configured for the selected Pool.
Option Type	Specifies the type of option associated with the option code configured for the selected pool. The possible values are as follows: • **Ascii**: The option type is a text string. • **Hex**: The option type is a hexadecimal number. • **IP Address**: The option type is an IP address.
ASCII Value	Shows the Option ASCII Value for the selected pool.
Hex Value	Shows the Option Hex Value for the selected pool.
IP Address Value	Shows the Option IP Address Value for the selected pool.
Delete Option Code	To delete an option code for the selected Pool, enter the option code in the folder and click **Delete**. This button is not visible to a user with read-only permission.

2.16.4 Reset Configuration

Use the Reset Configuration page to clear IP address bindings between that the DHCP server assigned to the client.

To access the Reset Configuration page, click **System > DHCP Server > Reset Configuration** in the navigation tree.

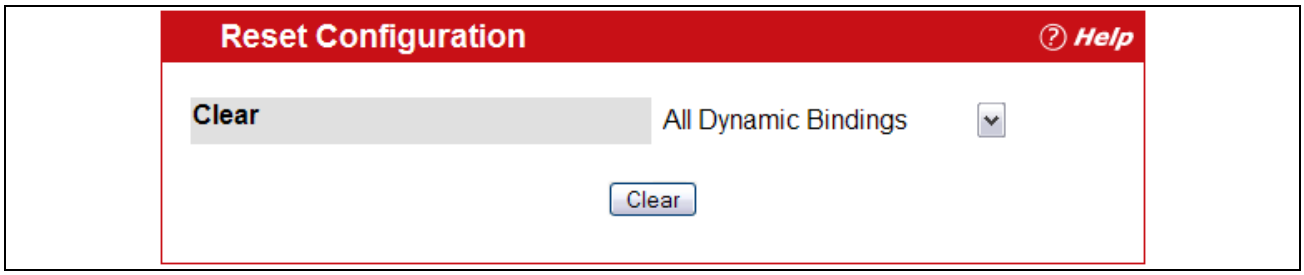

Figure 2-83: Reset Configuration

Table 2-76: Reset Configuration Fields

Field	Description
Clear	Specifies what to clear from the DHCP server database: • **All Dynamic Bindings**: Deletes all dynamic bindings from all address pools. • **Specific Dynamic Binding**: Deletes the specified binding. • **All Address Conflicts**: Deletes all address conflicts from the DHCP server database. • **Specific Address Conflict**: Deletes a specified conflicting address from the database.
Clear IP Address	If you select **Specific Dynamic Bindings** or **Specific Address Conflicts** from the **Clear** field, the screen refreshes and the **Clear IP Address** field appears. Enter the specific IP address to clear from the DHCP server.

After you select the bindings or conflicts to clear and, if necessary, enter the specific IP address, click **Clear** to remove the binding from the DHCP server.

2.16.5 Bindings Information

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

To access the Bindings Information page, click **System > DHCP Server > Bindings Information** in the navigation tree.

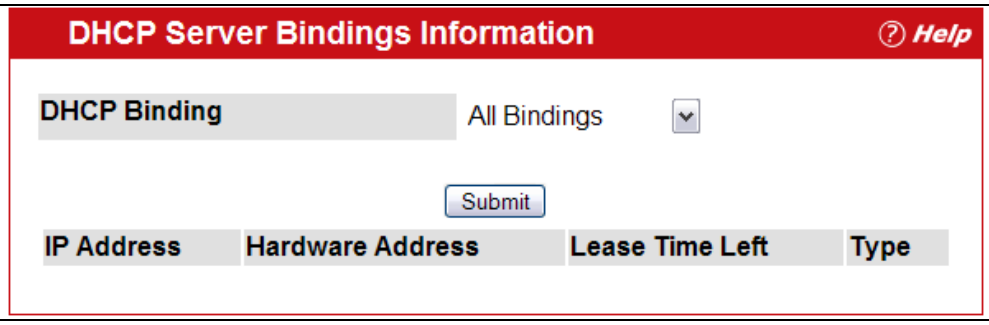

Figure 2-84: Bindings Information

Table 2-77: Bindings Information Fields

Field	Description
DHCP Binding	Select the bindings to display: • **All Bindings**: Show all bindings. • **Specific Binding**: Show a specific binding. When you select this option, the screen refreshes, and the **Binding IP Address** field appears.
Binding IP Address	Specify the IP address for which you want to view binding information. This field is only available if you select Specific Binding from the **DHCP Binding** field.

Table 2-77: Bindings Information Fields (Continued)

Field	Description
IP Address	Displays the client IP address.
Hardware Address	Displays the client MAC address.
Lease Time	Shows the remaining time left in the lease in Days, Hours and Minutes dd:hh:mm format.
Type	Shows the type of binding, which is dynamic or manual.

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

2.16.6 Server Statistics

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

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

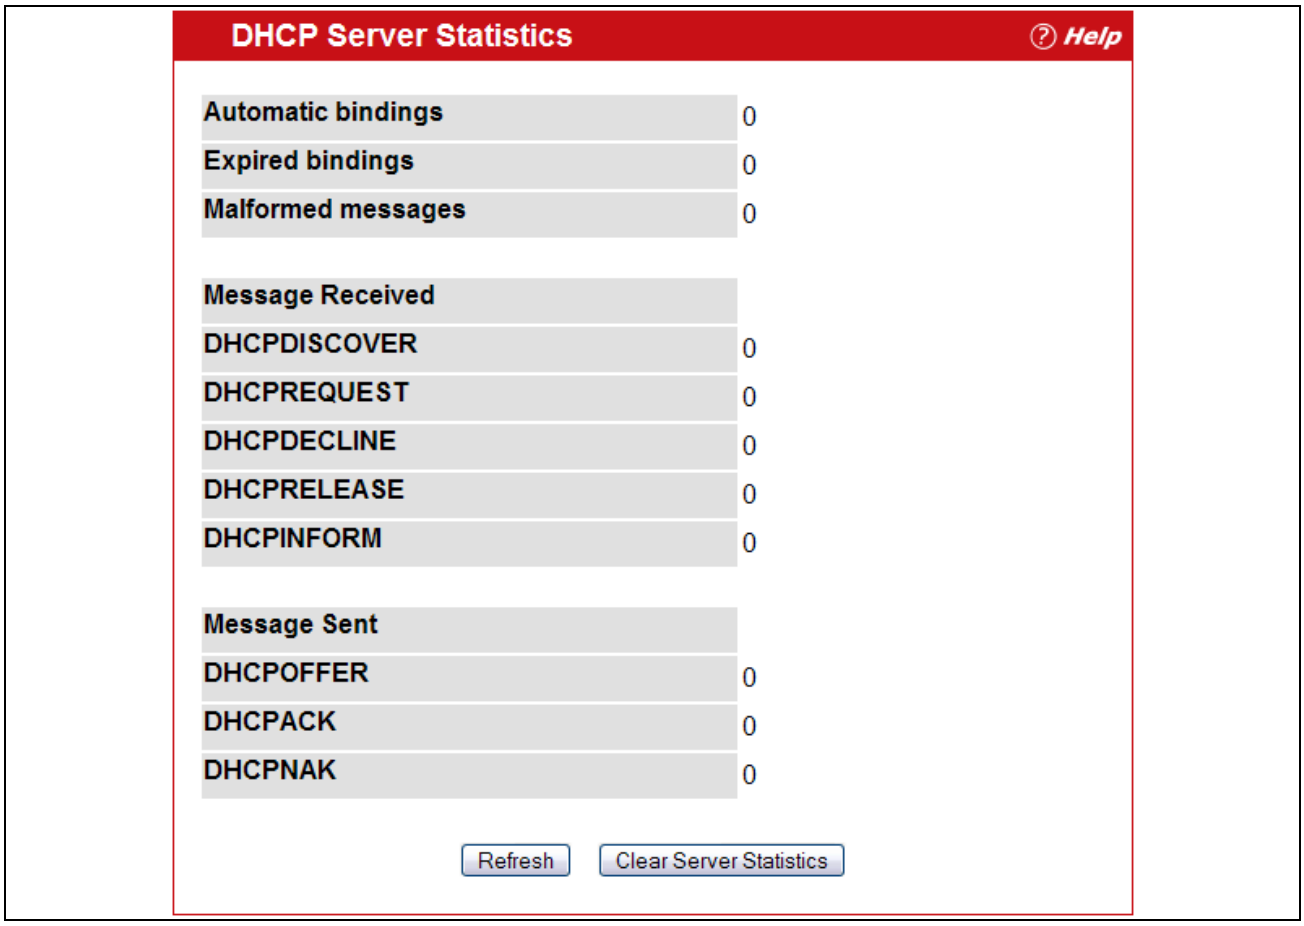

Figure 2-85: Server Statistics

Table 2-78: Server Statistics Fields

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

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

2.16.7 Conflicts Information

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

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

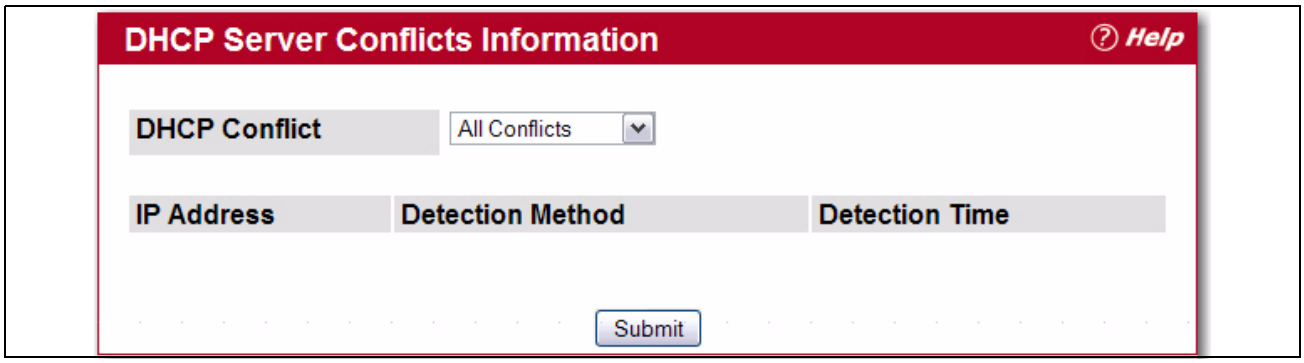

Figure 2-86: Conflicts Information

Table 2-79: Conflicts Information Fields

Field	Description
DHCP Conflicts	Select the DHCP conflicts to display: • **All Conflicts**: Show all conflicts. • **Specific Conflict**: Show a specific conflict. When you select this option, the screen refreshes, and the **Conflict IP Address** field appears.
Conflict IP Address	Specify the IP address for which you want to view conflict information. This field is only available if you select Specific Conflicts from the **DHCP Conflict** field.
IP Address	Displays the client IP address.
Detection Method	Specifies the manner in which the IP address of the hosts were found on the DHCP server.
Detection Time	Specifies the time when the conflict was detected in N days NNh:NNm:NNs format with respect to the system up time.

2.17 Configuring DNS

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

2.17.1 Global Configuration

Use this page to configure global DNS settings and to view DNS client status information.

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

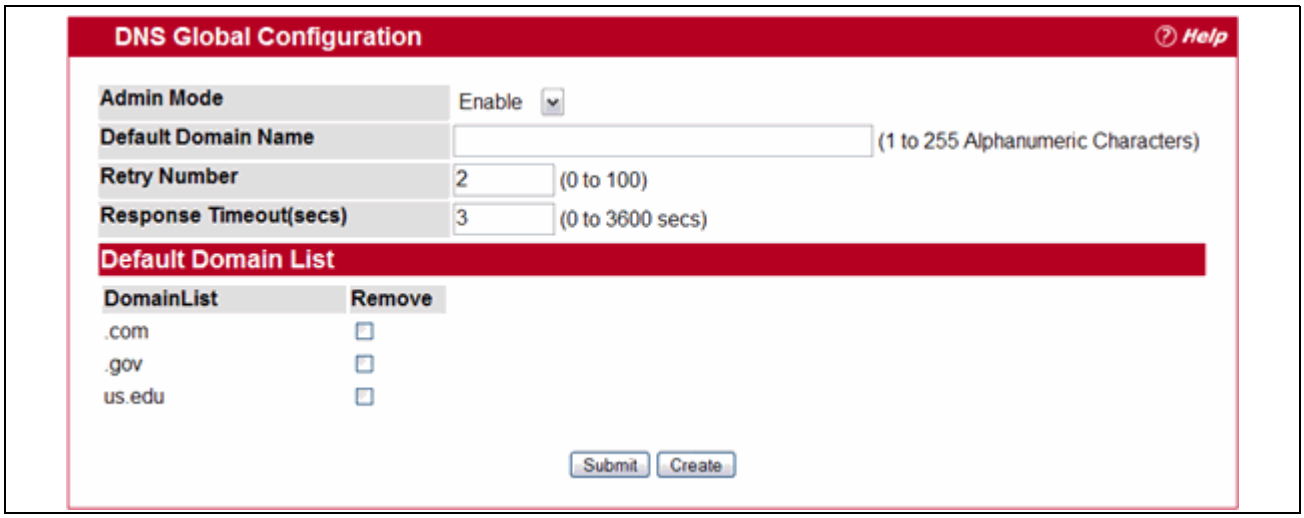

Figure 2-87: DNS Global Configuration

Table 2-80: DNS Global Configuration Fields

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

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

2.17.2 Server Configuration

Use this page to configure information about DNS servers that the router will use. The order in which you create them determines their precedence; i.e., DNS requests will go to the higher precedence server first. If that server is unavailable or does not respond in the configured response time, then the request goes to the server with the next highest precedence.

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

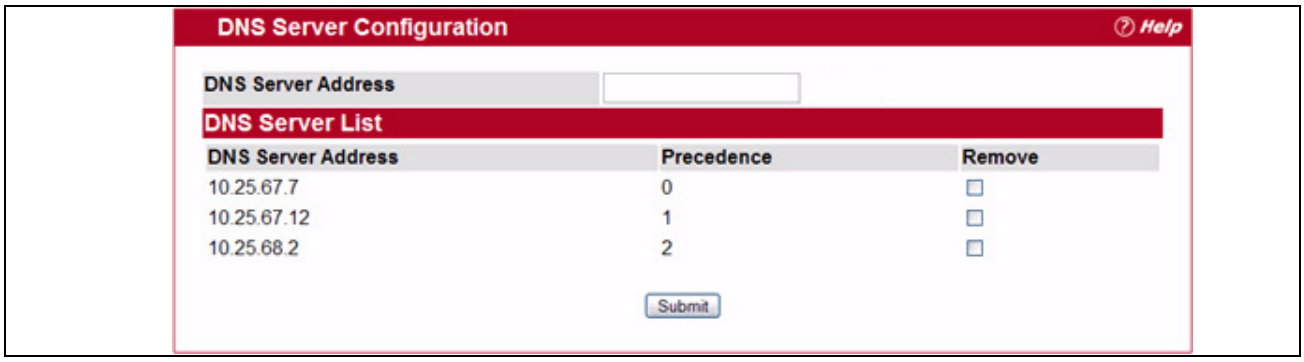

Figure 2-88: DNS Server Configuration

Table 2-81: DNS Server Configuration Fields

Field	Description
DNS Server Address	To add a new DNS server to the list, enter the DNS server IPv4 or IPv6 address in numeric notation.
Precedence	Shows the precedence value of the server that determines which server is contacted first; a lower number indicates has higher precedence.

- To create a new DNS server, enter an IP address in standard IPv4 or IPv6 dot notation in the **DNS Server Address** and click **Submit**. The server appears in the list below. The precedence is set in the order created.
- To change precedence, you must remove the server(s) by clicking the **Remove** box and then **Submit**, and add the server(s) in the preferred order.

2.17.3 DNS Host Name IP Mapping Configuration

Use this page to configure DNS host names for hosts on the network. The host names are associated with IPv4 or IPv6 addresses on the network, which are statically assigned to particular hosts.

To access this page, click **System > DNS > HostName IP Mapping Summary** in the navigation tree, then click the **Add Static Entry** button.

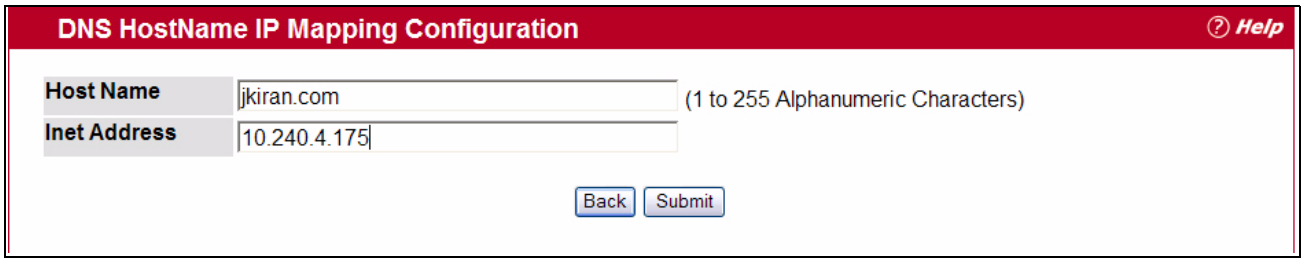

Figure 2-89: DNS Host Name Mapping Configuration

Table 2-82: DNS Host Name Mapping Configuration Fields

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

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

2.17.4 DNS Host Name IP Mapping Summary

Use this page to configure static and dynamic DNS host names for hosts on the network. The host names are associated with IPv4 or IPv6 addresses on the network, which are assigned to particular hosts.

To access this page, click **System > DNS > Host Name IP Mapping Summary** in the navigation tree.

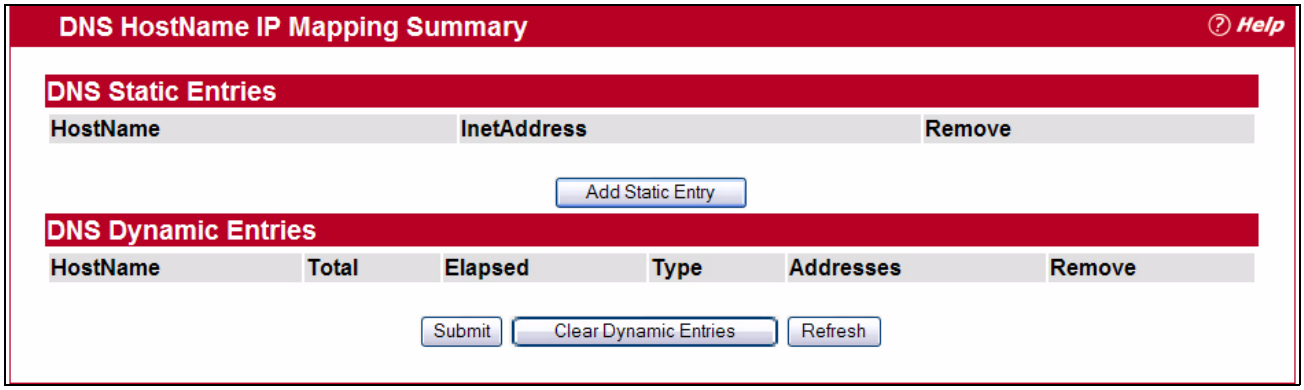

Figure 2-90: DNS Host Name IP Mapping Summary

Table 2-83: DNS Host Name IP Mapping Summary Fields

Field	Description
DNS Static Entries	
Host Name	The host name of the static entry.
Inet Address	The IP4 or IPv6 address of the static entry.
Remove	Select to remove a Host Name IP Mapping entry from the Host Name IP Mapping list.

Click **Add Static Entry** to load the Host Name IP Mapping Configuration page in order to configure the Host Name IP Mapping entries.

Table 2-84: DNS Host Name IP Mapping Configuration

Field	Description
DNS Dynamic Entries	
Host Name	The host name of the dynamic entry.
Total	The total time of the dynamic entry.
Elapsed	The elapsed time of the dynamic entry.
Type	The type of the dynamic entry.
Addresses	The IP4 or IPv6 address of the dynamic entry.
Remove	Select to remove a Host Name IP Mapping entry from the Host Name IP Mapping list.

- Click **Submit** to apply the new configuration and cause the change to take effect immediately. These changes will not be retained across a power cycle unless a Save is performed.

- Click **Clear Dynamic Entries** to remove all Host Name IP Mapping entries. A confirmation prompt will be displayed. Click the button to confirm removal and the Host Name IP Mapping dynamic entries are cleared.
- Click Refresh to refresh the page with the most current data from the switch.

2.18 Configuring SNTP Settings

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

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

The following is an example of stratums:

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

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

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

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

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

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

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

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

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

- If the servers have the same stratum, synchronization information is accepted from the SNTP server that responded first.

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

The SNTP folder contains links to view or configure the following features:

- SNTP Global Configuration
- SNTP Global Status
- SNTP Server Configuration
- SNTP Server Status

2.18.1 SNTP Global Configuration

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

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

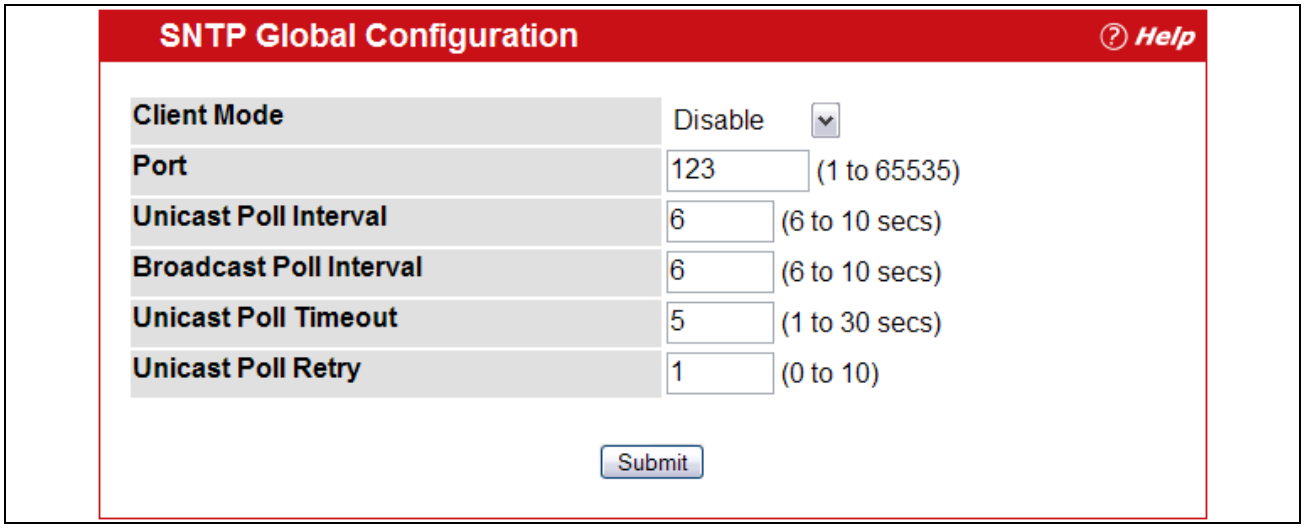

Figure 2-91: SNTP Global Configuration

Table 2-85: SNTP Global Configuration Fields

Field	Description
Client Mode	Use drop-down list specify the SNTP client mode, which is one of the following modes: • **Disable**: SNTP is not operational. No SNTP requests are sent from the client nor are any received SNTP messages processed. • **Unicast**: SNTP operates in a point to point fashion. A unicast client sends a request to a designated server at its unicast address and expects a reply from which it can determine the time and, optionally the round-trip delay and local clock offset relative to the server. • **Broadcast**: SNTP operates in the same manner as multicast mode but uses a local broadcast address instead of a multicast address. The broadcast address has a single subnet scope while a multicast address has Internet wide scope.
Port	Specifies the local UDP port to listen for responses/broadcasts. Allowed range is (1 to 65535). Default value is 123.
Unicast Poll Interval	Specifies the number of seconds between unicast poll requests expressed as a power of two when configured in unicast mode. Allowed range is (6 to 10). Default value is 6.
Broadcast Poll Interval	Specifies the number of seconds between broadcast poll requests expressed as a power of two when configured in broadcast mode. Broadcasts received prior to the expiry of this interval are discarded. Allowed range is (6 to 10). Default value is 6.
Unicast Poll Timeout	Specifies the number of seconds to wait for an SNTP response when configured in unicast mode. Allowed range is (1 to 30). Default value is 5.
Unicast Poll Retry	Specifies the number of times to retry a request to an SNTP server after the first time-out before attempting to use the next configured server when configured in unicast mode. Allowed range is (0 to 10). Default value is 1.

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

2.18.2 SNTP Global Status

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

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

SNTP Global Status Help

Version	4
Supported Mode	Unicast & Broadcast
Last Update Time	JAN 01 00:00:00 1970
Last Attempt Time	JAN 01 00:00:00 1970
Last Attempt Status	Other
Server IP Address	
Address Type	Unknown
Server Stratum	0 - Unspecified
Reference Clock Id	
Server Mode	Reserved
Unicast Server Max Entries	3
Unicast Server Current Entries	0
Broadcast Count	0

Figure 2-92: Global Status

Table 2-86: Global Status Fields

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

Table 2-86: Global Status Fields (Continued)

Field	Description
Address Type	Specifies the address type of the SNTP Server address for the last received valid packet.
Server Stratum	Specifies the claimed stratum of the server for the last received valid packet.
Reference Clock Id	Specifies the reference clock identifier of the server for the last received valid packet.
Server Mode	Specifies the mode of the server for the last received valid packet.
Unicast Sever Max Entries	Specifies the maximum number of unicast server entries that can be configured on this client.
Unicast Server Current Entries	Specifies the number of current valid unicast server entries configured for this client.
Broadcast Count	Specifies the number of unsolicited broadcast SNTP messages that have been received and processed by the SNTP client since last reboot.

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

2.18.3 SNTP Server Configuration

Use the SNTP Server Configuration page to view and modify information for adding and modifying Simple Network Time Protocol SNTP servers.

To display the SNTP Server Configuration page, click **System > SNTP > Server Configuration** in the navigation tree.

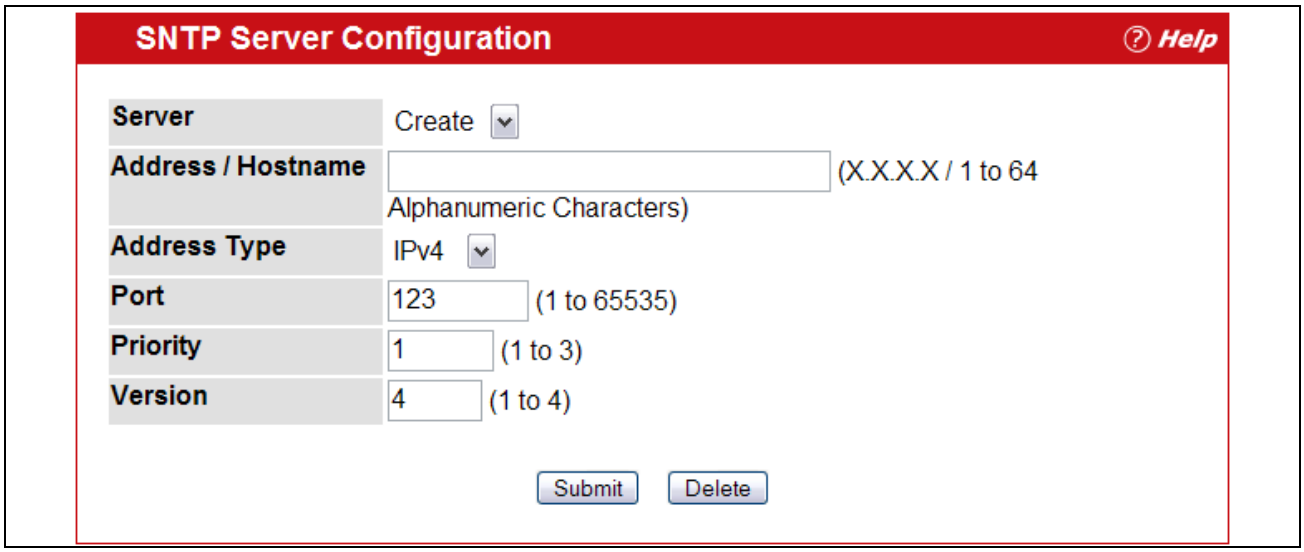

Figure 2-93: SNTP Server Configuration

Table 2-87: SNTP Server Configuration Fields

Field	Description
Server	Select the IP address of a user-defined SNTP server to view or modify information about an SNTP server, or select **Create** to configure a new SNTP server. You can define up to three SNTP servers.
Address / Hostname	Enter the IP address or the hostname of the SNTP server.
Address Type	Select **IPv4** if you entered an IPv4 address or **DNS** if you entered a hostname.
Port	Enter a port number from 1 to 65535. The default is 123.
Priority	Enter a priority from 1 to 3, with 1 being the highest priority. The router will attempt to use the highest priority server and, if it is not available, will use the next highest server.
Version	Enter the protocol version number.
Priority (1-3)	Specifies the priority of this server entry in determining the sequence of servers to which SNTP requests are sent. Values are 1 to 3, and the default is 1. Servers with lowest numbers have priority.

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

2.18.4 SNTP Server Status

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

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

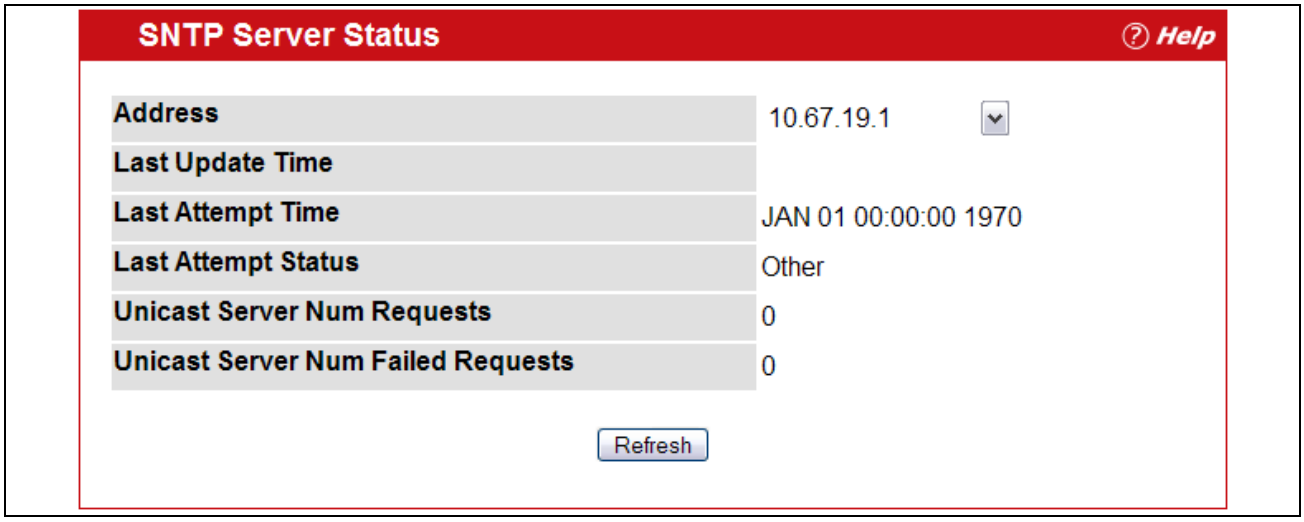

Figure 2-94: SNTP Server Status

Table 2-88: SNTP Server Status Fields

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

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

2.19 Configuring and Viewing ISDP Information

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

The following pages are accessible from this ISDP folder:

- Global Configuration
- Cache Table
- Interface Configuration
- Statistics

2.19.1 Global Configuration

From the ISDP **Global Configuration** page, you can configure the ISDP settings for the switch, such as the administrative mode.

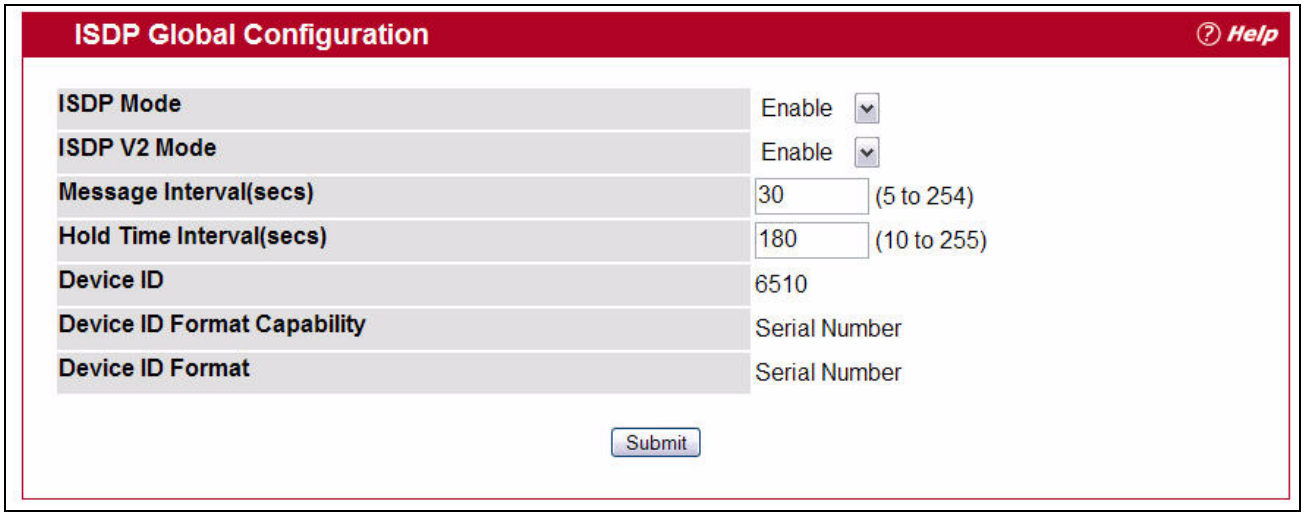

Figure 2-95: ISDP Global Configuration

The following table describes the fields available on the ISDP **Global Configuration** page.

Table 2-89: ISDP Global Configuration

Field	Description
ISDP Mode	Use this field to enable or disable the Industry Standard Discovery Protocol on the switch.
ISDP V2 Mode	Use this field to enable or disable the Industry Standard Discovery Protocol v2 on the switch.
Message Interval	Specifies the ISDP transmit interval. The range is (5–254). Default value is 30 seconds.
Holdtime Interval	The receiving device holds ISDP message during this time period. The range is (10–255). Default value is 180 seconds.
Device ID	The Device ID advertised by this device. The format of this Device ID is characterized by the value of Device ID Format object.
Device ID Format Capability	Indicates the Device ID format capability of the device. • serialNumber—Indicates that the device uses serial number as the format for its Device ID. • macAddress—Indicates that the device uses layer 2 MAC address as the format for its Device ID. • other—Indicates that the device uses its platform specific format as the format for its Device ID.
Device ID Format	Indicates the Device ID format of the device. • serialNumber—Indicates that the value is in the form of an ASCII string containing the device serial number. • macAddress—Indicates that the value is in the form of Layer 2 MAC address. • other—Indicates that the value is in the form of a platform specific ASCII string containing info that identifies the device. For example: ASCII string contains serialNumber appended/prepended with system name.

2.19.2 Cache Table

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

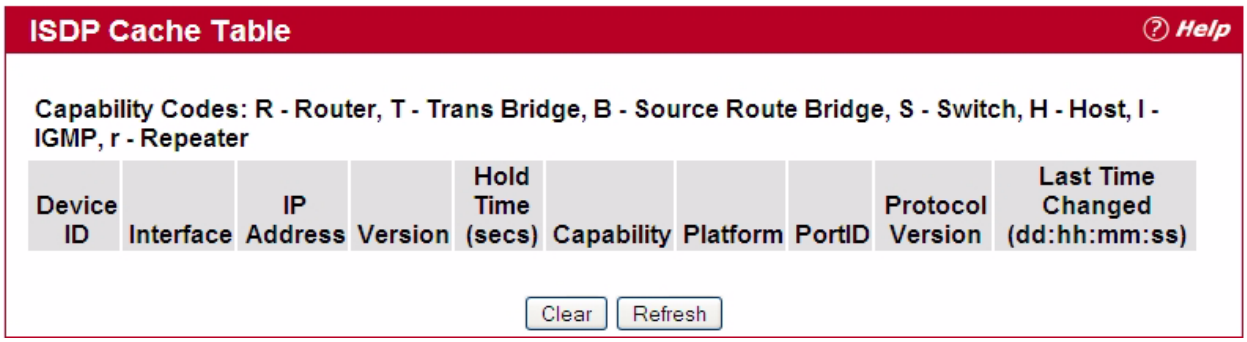

Figure 2-96: ISDP Cache Table

The following table describes the fields available on the ISDP **Cache Table** page.

Table 2-90: ISDP Cache Table

Field	Description
Device ID	Displays the string with Device ID which is reported in the most recent ISDP message.
Interface	Displays the interface that this neighbor is attached to.
IP Address	The (first) network-layer address that is reported in the Address TLV of the most recently received ISDP message.
Version	Displays the Version string for the neighbor.
Holdtime	Displays the ISDP holdtime for the neighbor.
Capability	Displays the ISDP Functional Capabilities for the neighbor.
Platform	Displays the ISDP Hardware Platform for the neighbor.
Port ID	Displays the ISDP port ID string for the neighbor.
Protocol Version	Displays the ISDP Protocol Version for the neighbor.
Last Time Changed	Displays when entry was last modified.

2.19.3 Interface Configuration

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

Note...

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

Figure 2-97: ISDP Interface Configuration

The following table describes the fields available on the ISDP **Interface Configuration** page.

Table 2-91: ISDP Interface Configuration

Field	Description
Slot/Port	Select the interface with the ISDP mode status to configure or view.
ISDP Mode	Use this field to enable or disable the Industry Standard Discovery Protocol on the selected interface.

2.19.4 Statistics

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

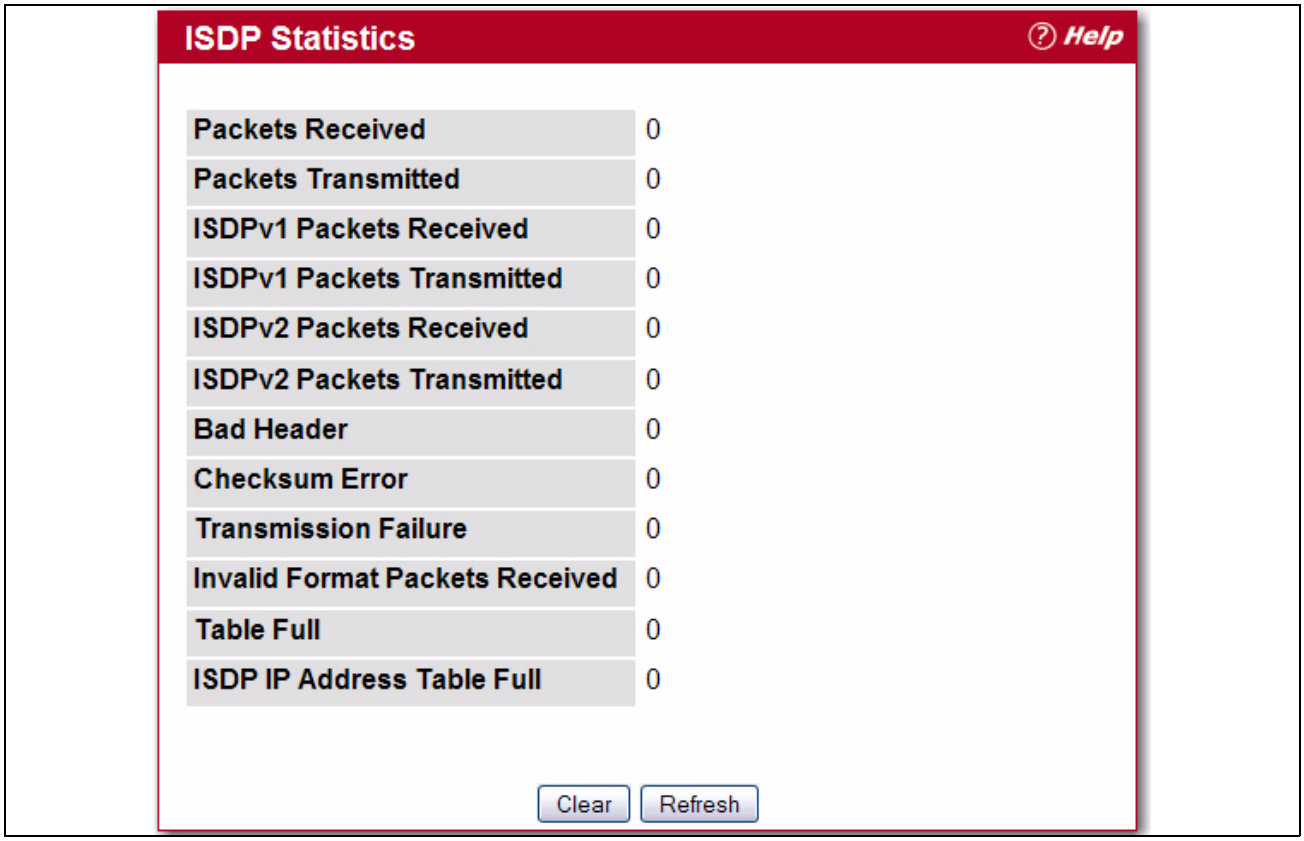

Figure 2-98: ISDP Statistics

The following table describes the fields available on the ISDP **Statistics** page.

Table 2-92: ISDP Statistics

Field	Description
ISDP Packets Received	Displays the number of all ISDP protocol data units (PDUs) received.
ISDP Packets Transmitted	Displays the number of all ISDP PDUs transmitted.
ISDPv1 Packets Received	Displays the number of v1 ISDP PDUs received.
ISDPv1 Packets Transmitted	Displays the number of v1 ISDP PDUs transmitted.
ISDPv2 Packets Received	Displays the number of v2 ISDP PDUs received.

Table 2-92: ISDP Statistics (Continued)

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

3 Configuring Switching Information

- Configuring DHCP Snooping
- Managing VLANs
- Double VLAN (DVLAN) Tunneling
- Configuring Protected Ports
- Managing IP Subnet-Based VLANs
- Managing MAC-Based VLANs
- Voice VLAN Configuration
- Creating MAC Filters
- Configuring GARP
- Configuring Dynamic ARP Inspection
- Configuring IGMP Snooping
- Configuring IGMP Snooping Queriers
- Configuring MLD Snooping
- Configuring MLD Snooping Queriers
- Creating Port Channels
- Viewing Multicast Forwarding Database Information
- Configuring Spanning Tree Protocol
- Mapping 802.1p Priority
- Configuring Port Security
- Managing LLDP

3.1 Configuring DHCP Snooping

DHCP snooping is a security feature that monitors DHCP messages between a DHCP client and DHCP servers to filter harmful DHCP messages and to build a bindings database of {MAC address, IP address, VLAN ID, port} tuples that are considered authorized. You can enable DHCP snooping globally and on specific VLANs, and configure ports within the VLAN to be trusted or untrusted. DHCP servers must be reached through trusted ports. DHCP snooping enforces the following security rules:

- DHCP packets from a DHCP server (DHCPOFFER, DHCPACK, DHCPNAK, DHCPRELEASEQUERY) are dropped if received on an untrusted port.
- DHCPRELEASE and DHCPDECLINE messages are dropped if destined for a MAC address in the snooping database, but the corresponding IP address in the snooping database is different than the interface where the message was received.
- On untrusted interfaces, the switch drops DHCP packets whose source MAC address does not match the client hardware address. This feature is a configurable option.

The hardware identifies all incoming DHCP packets on ports where DHCP snooping is enabled. DHCP snooping is enabled on a port if (a) DHCP snooping is enabled globally, and (b) the port is a member of a VLAN where DHCP snooping is enabled. On untrusted ports, the hardware traps all incoming DHCP packets to the CPU. On trusted ports, the hardware forwards client messages and copies server messages to the CPU so that DHCP snooping can learn the binding.

3.1.1 Global DHCP Snooping Configuration

To access the DHCP Snooping Configuration page, click **Switching > DHCP Snooping > Configuration** in the navigation tree.

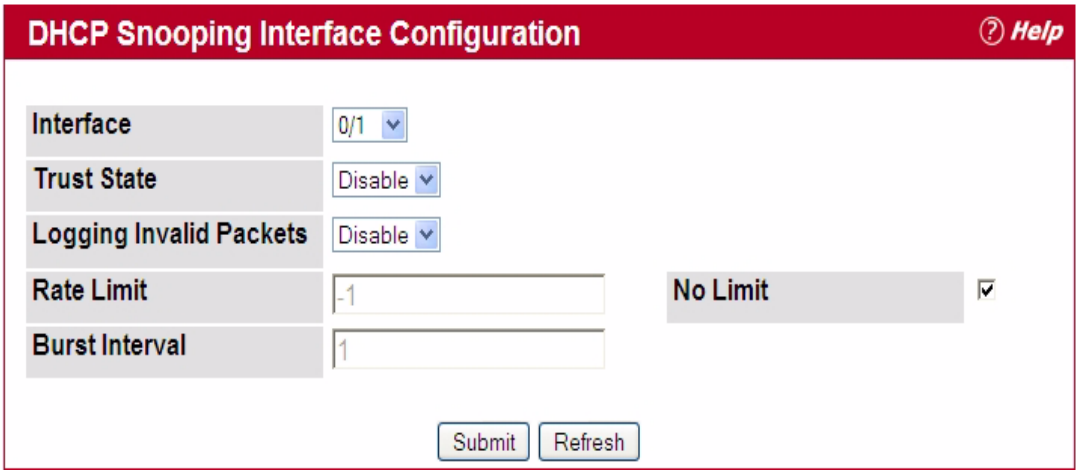

Figure 3-1: DHCP Snooping Configuration

Table 3-1: DHCP Snooping Configuration

Field	Description
DHCP Snooping Mode	Enables or disables the DHCP Snooping feature. The default is **Disable**.
MAC Address Validation	Enables or disables the validation of sender MAC Address for DHCP Snooping. The default is **Enable**.

Click **Submit** to apply the new configuration and cause the change to take effect. These changes will not be retained across a power cycle unless a **Save configuration** is performed.

3.1.2 DHCP Snooping VLAN Configuration

The DHCP snooping application does not forward server messages because they are forwarded in hardware.

DHCP snooping forwards valid DHCP client messages received on non-routing VLANs. The message is forwarded on all trusted interfaces in the VLAN.

DHCP snooping can be configured on switching VLANs and routing VLANs. When a DHCP packet is received on a routing VLAN, the DHCP snooping application applies its filtering rules and updates the bindings database. If a client message passes filtering rules, the message is placed into the software forwarding path, where it may be processed by the DHCP relay agent, the local DHCP server, or forwarded as an IP packet.

DHCP snooping is disabled globally and on all VLANs by default. Ports are untrusted by default.

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

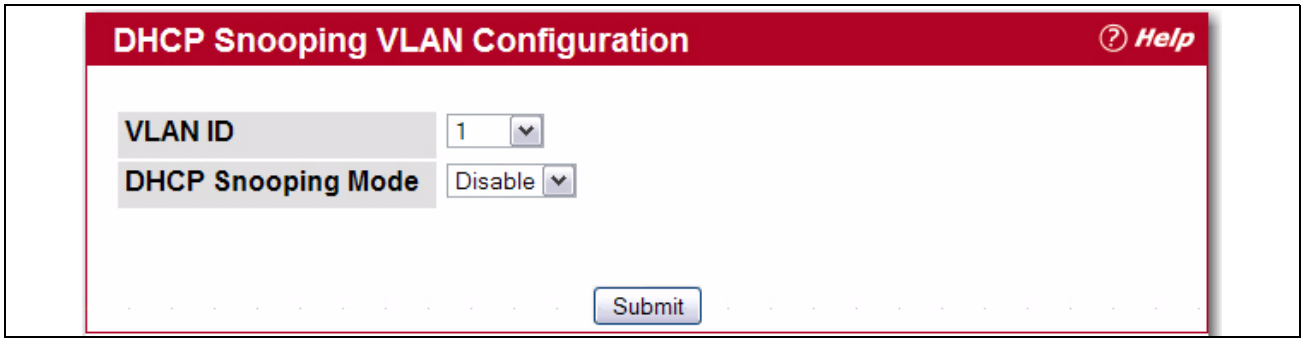

Figure 3-2: DHCP Snooping VLAN Configuration

Table 3-2: DHCP Snooping VLAN Configuration

Field	Description
VLAN ID	Select the VLAN for which information to be displayed or configured for the DHCP snooping application.
DHCP Snooping Mode	Enables or disables the DHCP snooping feature on the selected VLAN. The default is **Disable**.

Click **Submit** to apply the new configuration and cause the change to take effect. These changes will not be retained across a power cycle unless a **Save configuration** is performed.

3.1.3 DHCP Snooping Interface Configuration

The hardware rate limits DHCP packets sent to the CPU from untrusted interfaces to 64 kbps. There is no hardware rate limiting on trusted interfaces.

To prevent DHCP packets from being used as a DoS attack when DHCP snooping is enabled, the snooping application enforces a rate limit for DHCP packets received on untrusted interfaces. DHCP snooping monitors the receive rate on each interface separately. If the receive rate exceeds the configuration limit, DHCP snooping brings down the interface. You must do “no shutdown” on this interface to further work with that port. You can configure both the rate and the burst interval.

The DHCP snooping application processes incoming DHCP messages. For DHCPRELEASE and DHCPDECLINE messages, the application compares the receive interface and VLAN with the client's interface and VLAN in the binding database. If the interfaces do not match, the application logs the event and drops the message. For valid client messages, DHCP snooping compares the source MAC address to the DHCP client hardware address. Where there is a mismatch, DHCP snooping logs and drops the packet. You can disable this feature using the DHCP Snooping Interface Configuration page, shown in Figure 3-3 below, or by using the `no ip dhcp snooping verify mac-address` command. DHCP snooping forwards valid client messages on trusted members within the VLAN. If DHCP relay and/or DHCP server co-exist with the DHCP snooping, the DHCP client message will be sent to the DHCP relay and/or DHCP server to process further.

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

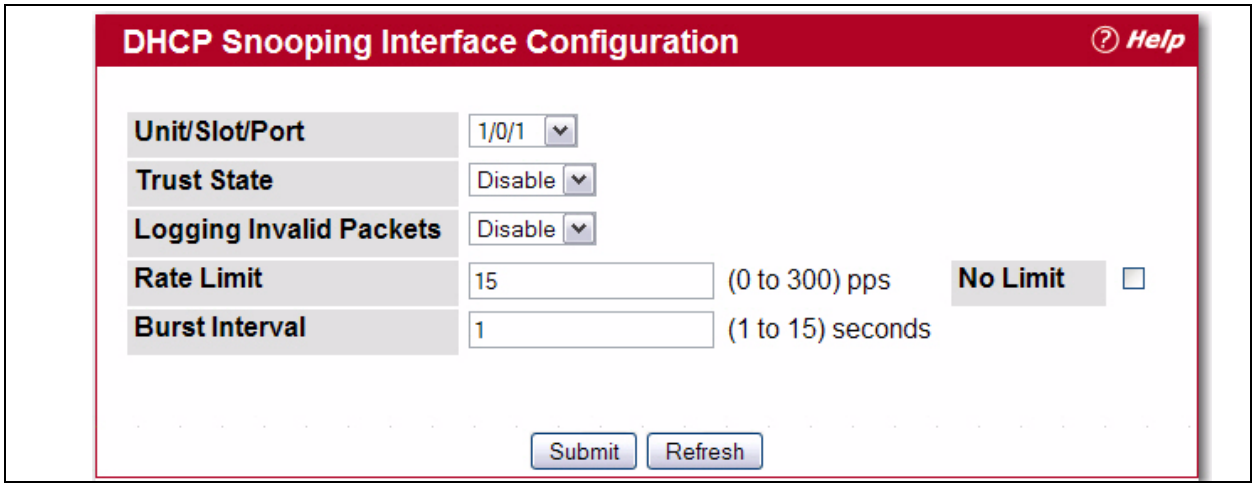

Figure 3-3: DHCP Snooping Interface Configuration

Table 3-3: DHCP Snooping Interface Configuration

Field	Description
Slot/Port	Select the interface for which data is to be displayed or configured.
Trust State	If it is enabled, the DHCP snooping application considers the port as trusted. The default is **Disable**.
Logging Invalid Packets	If it is enabled, the DHCP snooping application logs invalid packets on this interface. The default is **Disable**.
Rate Limit	Specifies the rate limit value for DHCP snooping purposes. If the incoming rate of DHCP packets exceeds the value of this object for consecutively burst interval seconds, the port will be shutdown. If this value is **None**, there is no limit. The default is 15 packets per second (pps). The Rate Limit range is 0 to 300.
Burst Interval	Specifies the burst interval value for rate limiting purposes on this interface. If the rate limit is **None**, the burst interval has no meaning and displays it as "N/A". The default is **1 second**. The Burst Interval range is 1 to 15.

Click **Submit** to apply the new configuration and cause the change to take effect. These changes will not be retained across a power cycle unless a **Save configuration** is performed.

3.1.4 DHCP Snooping Binding Configuration

The DHCP snooping application uses DHCP messages to build and maintain the binding's database. The binding's database only includes data for clients on untrusted ports. DHCP snooping creates a *tentative binding* from DHCP DISCOVER and REQUEST messages. Tentative bindings tie a client to a port (the port where the DHCP client message was received). Tentative bindings are completed when DHCP snooping learns the client's IP address from a DHCP ACK message on a trusted port. DHCP snooping removes bindings in response to DECLINE, RELEASE, and NACK messages. The DHCP snooping application ignores the ACK messages as a reply to the DHCP Inform messages received on trusted ports. You can also enter static bindings into the binding database.

The DHCP binding database is persisted on a configured external server or locally in flash, depending on the user configuration. A row wise checksum is placed in the text file that is going to be stored in the remote configured server. On reloading, the switch reads the configured binding file to build the DHCP snooping database. When the switch starts and the calculated checksum value equals the stored checksum, the switch reads entries from the binding file and populates the binding database. A checksum failure or a connection problem to the external configured server will cause the switch to loose the bindings and will cause a host's data loss if IP Source Guard (IPSG) and/or DAI is enabled.

When a switch learns of new bindings or when it loses bindings, the switch immediately updates the entries in the database. The switch also updates the entries in the binding file. The frequency at which the file is updated is based on a configurable delay, and the updates are batched.

If the absolute lease time of the snooping database entry expires, that entry is removed. You should take care of the system time to be consistent across the reboots. Otherwise, the snooping entries will not expire properly. If a host sends a DHCP release while the switch is rebooting, when the switch receives the DHCP discovery or request, the client's binding goes to the tentative binding as shown in Figure 3-4 on page 154.

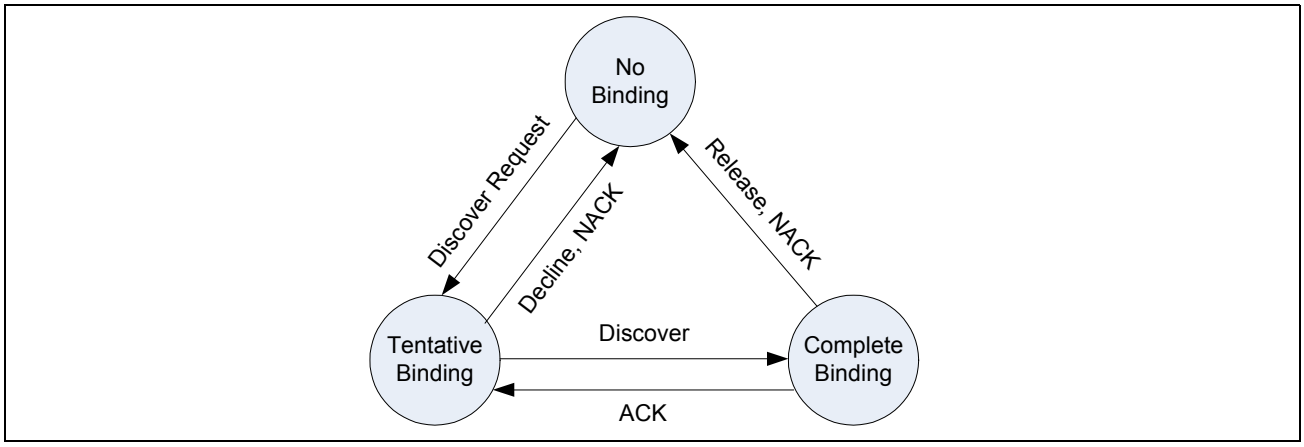

Figure 3-4: States of Client Binding

To access the DHCP Snooping Static Binding Configuration page, click **Switching > DHCP Snooping > Binding Configuration** in the navigation tree.

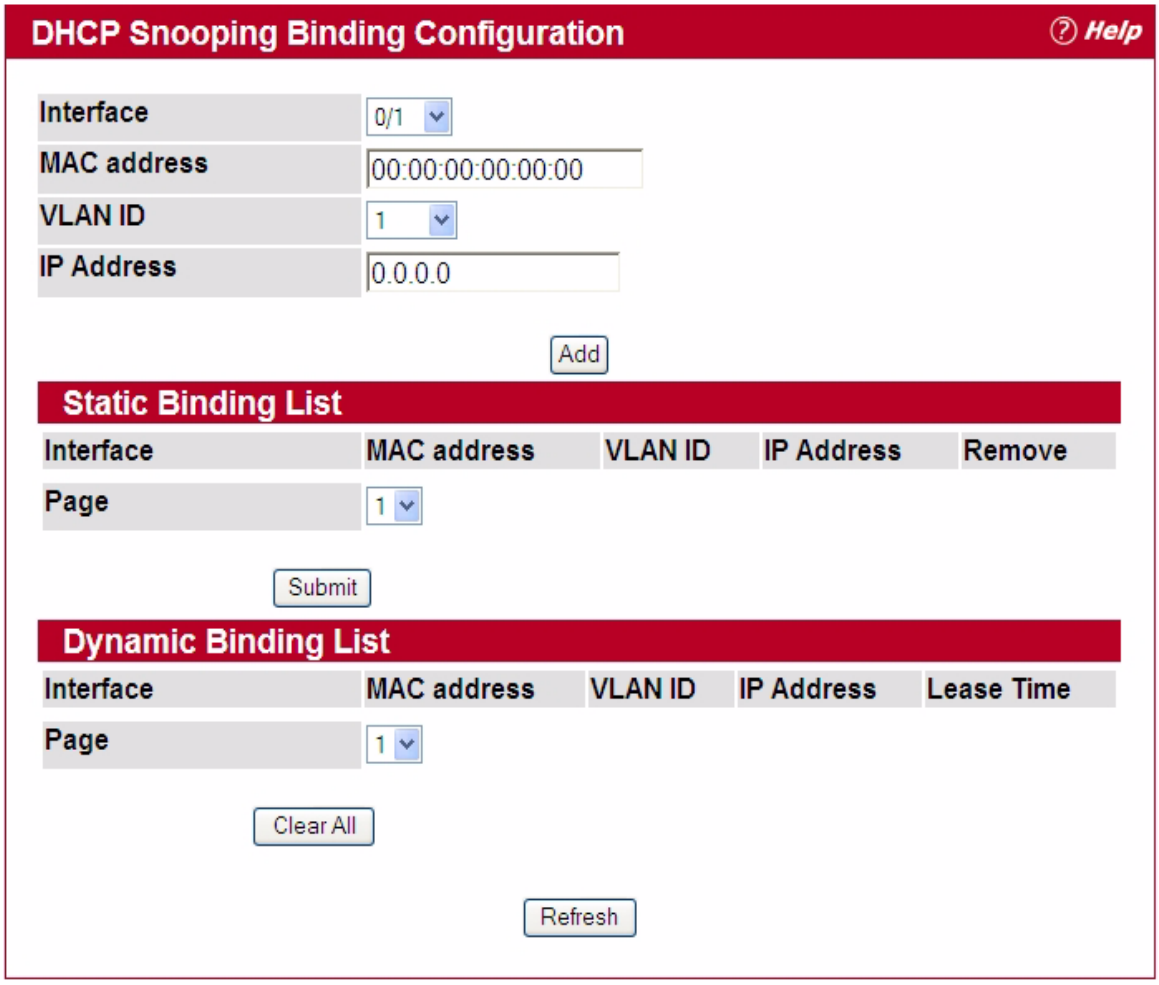

Figure 3-5: DHCP Snooping Binding Configuration

Table 3-4: DHCP Snooping Static Binding Configuration

Field	Description
Slot/Port	Select the interface to add a binding into the DHCP snooping database.
MAC Address	Specify the MAC address for the binding to be added. This is the Key to the binding database.
VLAN ID	Select the VLAN from the list for the binding rule. The range of the VLAN ID is 1 to 3965.
IP Address	Specify a valid IP address for the binding rule.

The DHCP snooping static binding list lists all the DHCP snooping static binding entries page by page. For example, **Page 1** displays the first 15 available static entries. **Page 2** displays the next 15 available static entries.

Table 3-5: DHCP Snooping Static Binding List

Field	Description
Slot/Port	Displays the interface.
MAC Address	Displays the MAC address.
VLAN ID	Displays the VLAN ID .
IP Address	Displays the IP address.
Remove	Select this to remove the particular binding entry.
Page	Lists the number of pages the static binding entries occupy. Select the Page Number from this list to display the particular Page entries.

The DHCP snooping dynamic binding list lists all the DHCP snooping dynamic binding entries page by page. For example, **Page 1** displays the first 15 available dynamic entries. **Page 2** displays the next 15 available dynamic entries.

Table 3-6: DHCP Snooping Dynamic Binding List

Field	Description
Slot/Port	Displays the interface.
MAC Address	Displays the MAC address.
VLAN ID	Displays the VLAN ID.
IP Address	Displays the IP address.
Lease Time	Displays the remaining Lease time for the dynamic entries.
Page	Lists the number of pages the static binding entries occupy. Select the Page Number from this list to display the particular Page entries.

- Click **Add** to add a DHCP snooping binding entry into the database.
- Click **Submit** to apply the new configuration and cause the change to take effect. These changes will not be retained across a power cycle unless a Save configuration is performed.
- Click **Clear All** to delete all DHCP snooping binding entries.
- Click **Refresh** to refresh the page with the most current data from the switch.

3.1.5 DHCP Snooping Persistent Configuration

Use the DHCP Snooping Persistent Configuration page to configure the persistent location of the DHCP snooping database. This location can be local or remote on a given IP machine.

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

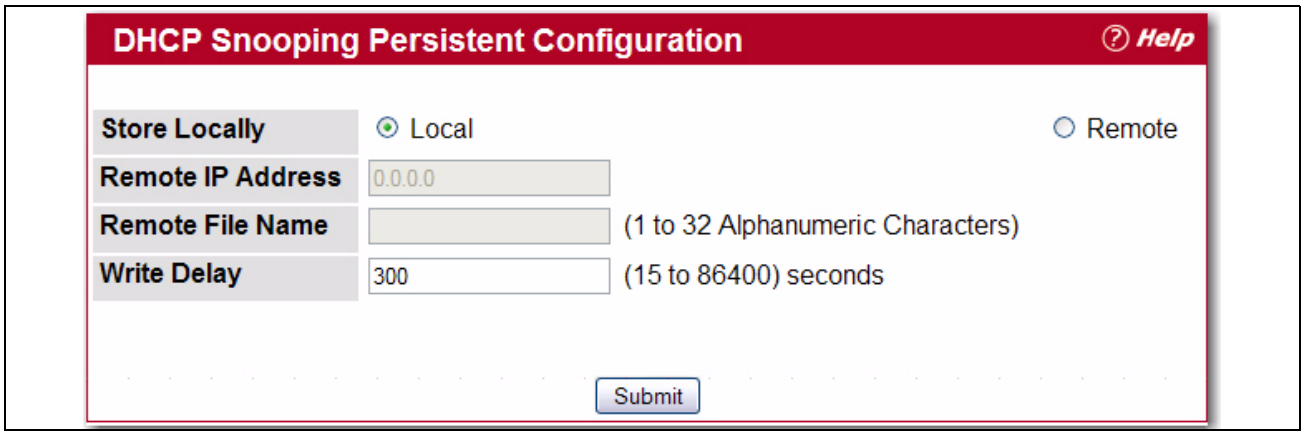

Figure 3-6: DHCP Snooping Persistent Configuration

Table 3-7: DHCP Snooping Persistent Configuration

Field	Description
Store Locally	• **Local**: Select the **Local** check box to store the DHCP binding database in the flash memory on the switch. • **Remote**: Check the **Remote** check box to store the DHCP binding database on a remote server.
Remote IP Address	Enter the Remote IP address on which the snooping database will be stored when the Remote check box is selected.
Remote File Name	Enter the Remote filename to store the database when the Remote check box is selected.
Write Delay	Enter the maximum write time to write the database into local or remote. The write delay range is 15 to 86400 seconds.

Click **Submit** to apply the new configuration and cause the change to take effect. These changes will not be retained across a power cycle unless a Save configuration is performed.

3.1.6 DHCP Snooping Statistics

The DHCP Snooping Statistics page displays DHCP snooping interface statistics.

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

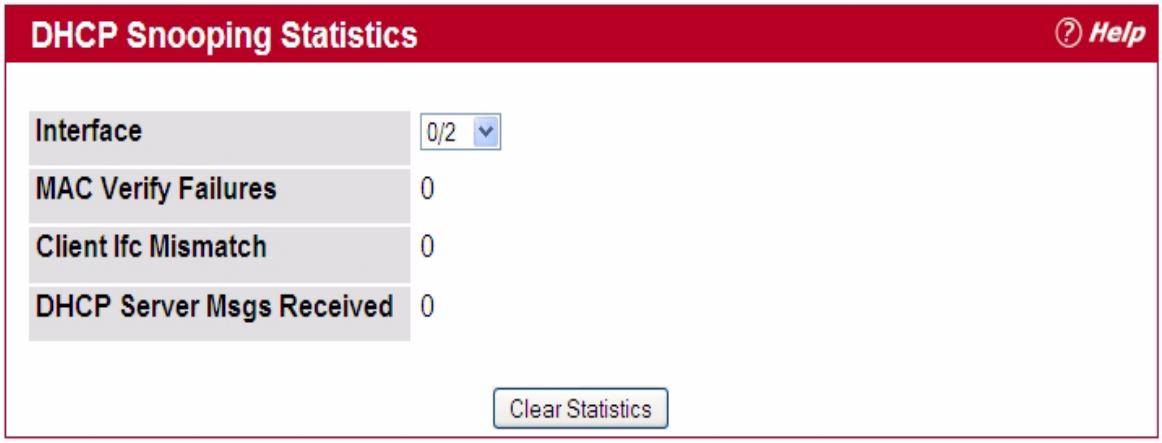

Figure 3-7: DHCP Snooping Statistics

Table 3-8: DHCP Snooping Statistics

Field	Description
Slot/Port	Select the untrusted and snooping-enabled interface for which statistics are to be displayed.
MAC Verify Failures	The number of packets that were dropped by DHCP snooping because there is no matching DHCP snooping binding entry found.
Client Ifc Mismatch	The number of DHCP messages that are dropped based on the source MAC address and client hardware address verification.
DHCP Server Msgs Received	The number of server messages that are dropped on an untrusted port.

Click the **Clear Stats** to clear all interface statistics.

3.1.7 Configuring DHCP L2 Relay

When a DHCP client and server are in the same IP subnet, they can directly connect to exchange IP address requests and replies. However, having a DHCP server on each subnet can be expensive in and is often impractical. Alternatively, network infrastructure devices can be used to relay packets between a DHCP client and server on different subnets. Such a device, a Layer 3 Relay agent, is generally a router that has IP interfaces on both the client and server subnets and can route between them. However, in Layer 2 switched networks, there may be one or more infrastructure devices (for example, a switch) between the client and the L3 Relay agent/DHCP server. In this instance, some of the client device information required by the L3 Relay agent may not be visible to it. In this case, an L2 Relay agent can be used to add the information that the L3 Relay Agent and DHCP server need to perform their roles in address and configuration and assignment.

Before it relays DHCP requests from clients, the switch can add a Circuit ID and a Remote ID. These provide information about the circuit and port number connected to the client. This information is added as suboptions in the DHCP Option 82 packets (see sections 3.1 and 3.2 of RFC3046). The switch removes this option from packets that it relays from L3 Relay agents/DHCP servers to clients.

These sub-options may be used by the DHCP server to affect how it treats the client, and also may be used by the relay agent to limit broadcast replies to the specific circuit or attachment point of the client.

The Switching > DHCP Snooping > DHCP L2 Relay folder provides access to the following pages:

- DHCP L2 Relay Global Configuration
- DHCP L2 Relay Interface Configuration
- DHCP L2 Relay VLAN Configuration
- DHCP L2 Relay Interface Statistics

3.1.7.1 DHCP L2 Relay Global Configuration

Use this page to enable or disable the switch to act as a DHCP L2 relay agent. This functionality must also be enabled on each port you want this service to operate on (see 3.1.7.2 DHCP L2 Relay Interface Configuration). The switch can also be configured to relay requests only when the VLAN of the requesting client corresponds to a service provider's VLAN ID that has been enabled with the L2 DHCP relay functionality (see "DHCHP L2 Relay VLAN Configuration" on page 115).

To access this page, click **Switching > DHCP Snooping > DHCP L2 Relay > Global Configuration**.

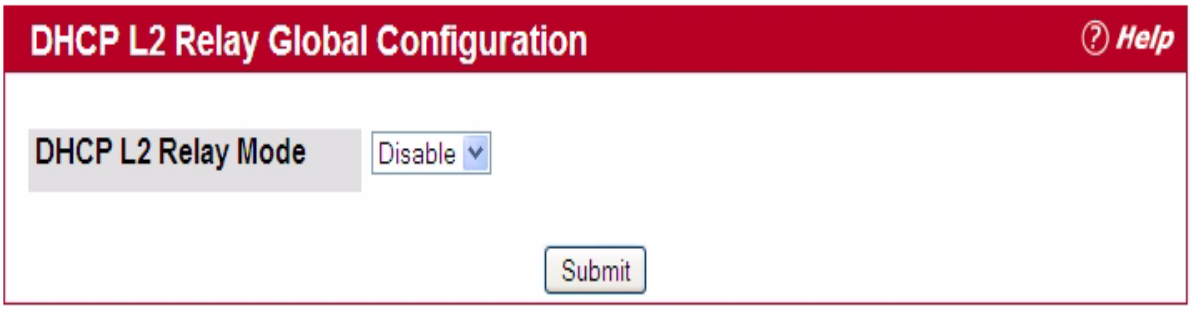

Figure 3-8: DHCP L2 Relay Global Configuration

If you enable or disable this feature, click Submit to apply the changes to system.

3.1.7.2 DHCP L2 Relay Interface Configuration

Use this page to enable L2 DHCP relay on individual ports. Note that L2 DHCP relay must also be enabled globally on the switch. To access this page, click **Switching > DHCP Snooping > DHCP L2 Relay > Interface Configuration**.

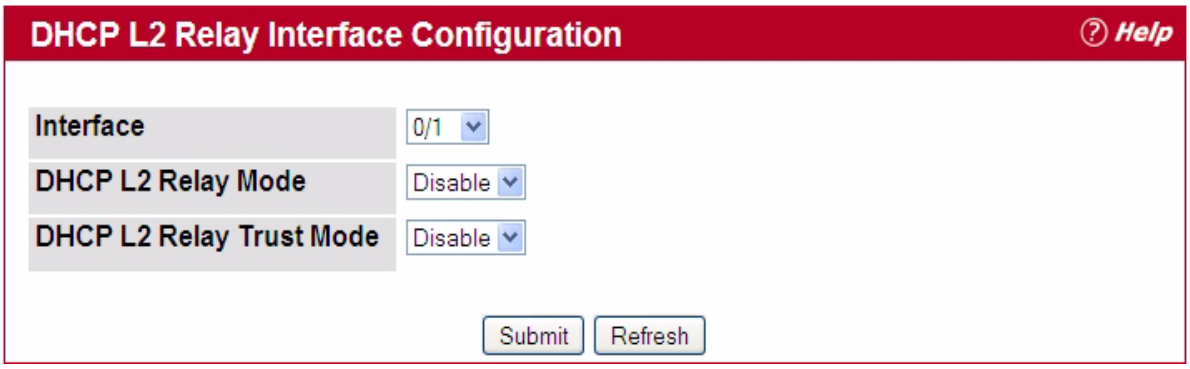

Figure 3-9: DHCP L2 Relay Interface Configuration

Table 3-9: DHCP L2 Relay Interface Configuration

Field	Description
Slot/Port	Select the slot/port to configure this feature on.
DHCP L2 Relay Mode	Enable or disable L2 Relay mode on the selected interface.
DHCP L2 Relay Trust Mode	Enable or disable L2 Relay Trust Mode on the selected interface. Trusted interfaces usually connect to other agents or servers participating in the DHCP interaction (e.g. other L2 or L3 Relay Agents or Servers). When enabled in Trust Mode, the interface always expects to receive DHCP packets that include Option 82 information. If Option 82 information is not included, these packets are discarded. Untrusted interfaces are generally connected to clients. DHCP packets arriving on an untrusted interface are never expected to carry Option 82 and are discarded if they do.

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

3.1.7.3 DHCP L2 Relay VLAN Configuration

You can enable L2 DHCP relay on a particular VLAN. The VLAN is identified by a service VLAN ID (S-VID), which a service provider uses to identify a customer's traffic while traversing the provider network to multiple remote sites. The switch uses the VLAN membership of the switch port client (the customer VLAN ID, or C-VID) to perform a lookup a corresponding S-VID.

If the S-VID is enabled for DHCP L2 Relay, the packet can be forwarded. If the C-VID does not correspond to an S-VID that is enabled for DHCP L2 relay, the switch will not relay the DHCP request packet.

To access this page, click **Switching > DHCP Snooping > DHCP L2 Relay > VLAN Configuration**.

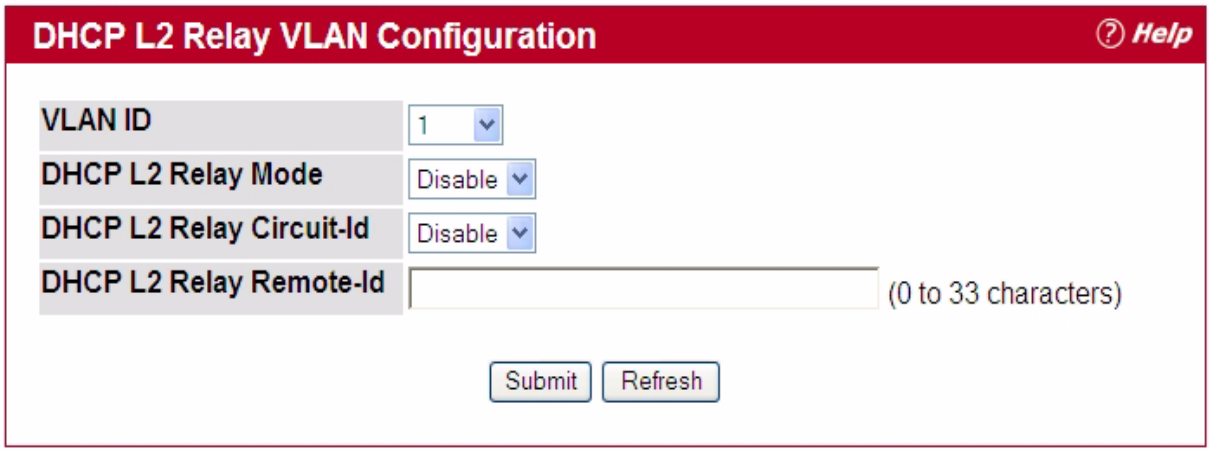

Figure 3-10: DHCP L2 Relay VLAN Configuration

Table 3-10: DHCP L2 Relay VLAN Configuration

Field	Description
VLAN ID	Select a VLAN ID from the list for configuration. This is an S-VID (as indicated by the service provider) that identifies a VLAN that is authorized to relay DHCP packets through the provider network.
DHCP L2 Relay Mode	Enable or disable the selected VLAN for DHCP L2 relay services.
DHCP L2 Relay Circuit-Id	When enabled, if a client sends a DHCP request to the switch and the client is in a VLAN that corresponds to the selected S-VID, the switch adds the client's interface number to the Circuit ID sub-option of Option 82 in the DHCP request packet. This enables the switch to reduce the broadcast domain to which the server replies are switched when the broadcast bit is set for DHCP packets. When this bit is set, the server is required to echo the Option-82 in replies. Since the circuit-id field contains the client interface number, the L2 relay agent can forward the response to the requesting interface only, rather to all ports in the VLAN).
DHCP L2 Relay Remote-Id	When a string is entered here, if a client sends a DHCP request to the switch and the client is in a VLAN that corresponds to the selected S-VID, the switch adds the string to the Remote-ID sub-option of Option 82 in the DHCP request packet. This sub-option can be used by the server for parameter assignment. The content of this option is vendor-specific.

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

3.1.7.4 DHCP L2 Relay Interface Statistics

Use this page to display statistics on L2 DHCP Relay requests received on a selected port. To access this page, click **Switching > DHCP Snooping > DHCP L2 Relay > Interface Statistics**.

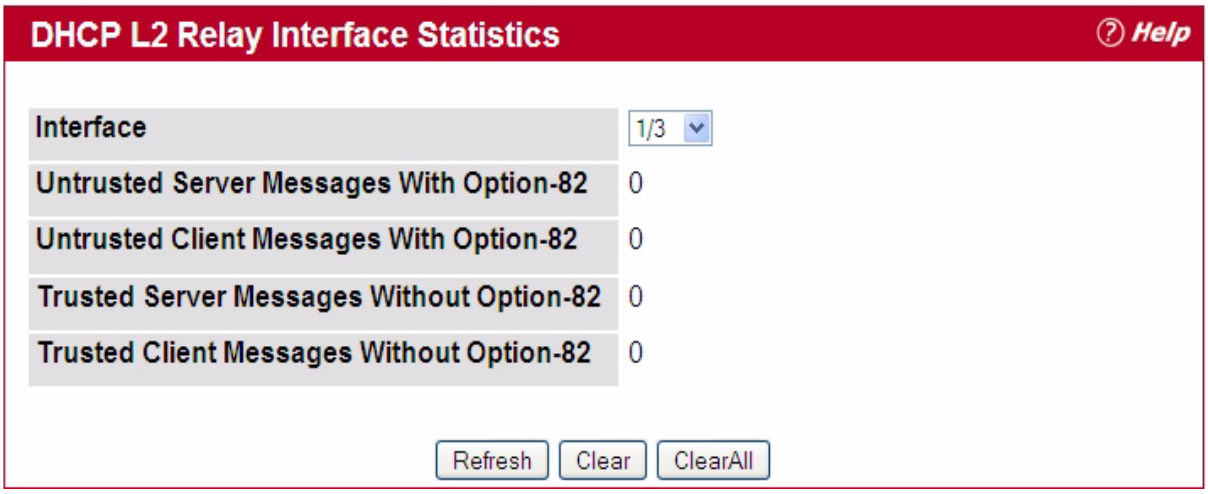

Figure 3-11: DHCP L2 Relay Interface Statistics

Table 3-11: DHCP L2 Relay Interface Statistics

Field	Description
Slot/Port	Select the slot/port to configure this feature on.
Untrusted Server Msgs With Option—82	If the selected interface is configured in untrusted mode, this field shows the number of messages received on the interface from a DHCP server that contained Option 82 data.These messages are dropped.
Untrusted Client Msgs With Option—82	If the selected interface is configured in untrusted mode, this field shows the number of messages received on the interface from a DHCP client that contained Option 82 data.These messages are dropped.
Trusted Server Msgs Without Option—82	If the selected interface is configured in trusted mode, this field shows the number of messages received on the interface from a DHCP server that did not contain Option 82 data. These messages are dropped.
Trusted Client Msgs Without Option—82	If the selected interface is configured in trusted mode, this field shows the number of messages received on the interface from a DHCP client that did not contain Option 82 data. These messages are dropped.

Click **Refresh** to redisplay the page with the latest information from the switch.

Click **Clear** to set statistics for this port to their initial values.

Click **Clear All** to set statistics for all ports to their initial values.

REVIEW QUESTION: Merlion also has DHCP L2 Relay Subscription Configuration & Subscription Summary pages. I don't see these pages on the ENT switch. Are they available on the SMB switch?

3.2 Managing VLANs

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

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

Each VLAN in a network has an associated VLAN ID, which appears in the IEEE 802.1Q tag in the Layer 2 header of packets transmitted on a VLAN. An end station may omit the tag, or the VLAN portion of the tag, in which case the first switch port to receive the packet may either reject it or insert a tag using its default VLAN ID. A given port may handle traffic for more than one VLAN, but it can only support one default VLAN ID.

The VLAN folder contains links to the following features:

- VLAN Configuration
- VLAN Status
- VLAN Port Configuration
- VLAN Port Summary
- Reset VLAN Configuration
- VLAN Internal Usage Configuration

3.2.1 VLAN Configuration

Use the VLAN Configuration page to define VLAN groups stored in the VLAN membership table. Your switch supports up to 3965 VLANs. VLAN 1 is the default VLAN of which all ports are members.

To display the VLAN Configuration page, click **Switching > VLAN > Configuration** in the navigation tree.

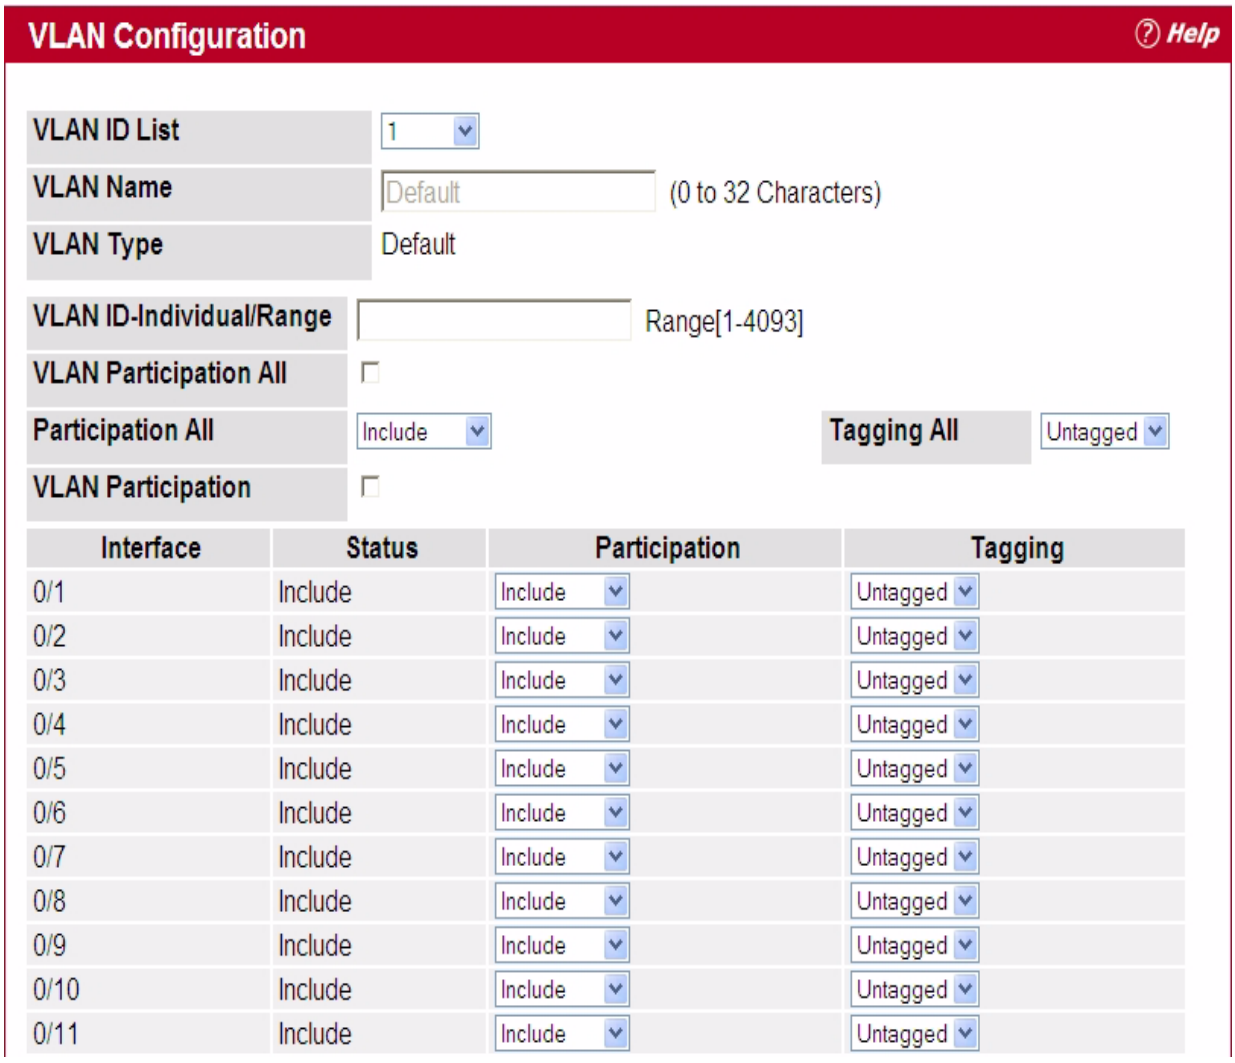

Figure 3-12: VLAN Configuration

Table 3-12: VLAN Configuration Fields

Field	Description
VLAN ID and Name	You can use this screen to reconfigure an existing VLAN, or to create a new one. Use this pulldown menu to select one of the existing VLANs, or select **Create** to add a new one.
VLAN ID	Specify the VLAN Identifier for the new VLAN. (You can only enter data in this field when you are creating a new VLAN.) The range of the VLAN ID is (1 to 3965).
VLAN Name	Use this optional field to specify a name for the VLAN. It can be up to 32 alphanumeric characters long, including blanks. The default is blank. VLAN ID 1 is always named “Default.”
VLAN Type	This field identifies the type of the VLAN you are configuring. You cannot change the type of the default VLAN (VLAN ID = 1): it is always type “Default.” When you create a VLAN, using this screen, its type will always be “Static.” A VLAN that is created by GVRP registration initially has a type of “Dynamic.” You can use this pulldown menu to change its type to “Static.”

Table 3-12: VLAN Configuration Fields (Continued)

Field	Description
Slot/Port	Indicates which port is associated with the fields on this line.
Status	Indicates the current value of the participation parameter for the port.
Participation	Use this field to specify whether a port will participate in this VLAN. The factory default is "Autodetect." The possible values are: • **Include**: This port is always a member of this VLAN. This is equivalent to registration fixed in the IEEE 802.1Q standard. • **Exclude**: This port is never a member of this VLAN. This is equivalent to registration forbidden in the IEEE 802.1Q standard. • **Autodetect**: Specifies that port may be dynamically registered in this VLAN via GVRP. The port will not participate in this VLAN unless it receives a GVRP request. This is equivalent to registration normal in the IEEE 802.1Q standard.
Tagging	Select the tagging behavior for this port in this VLAN. The factory default is "Untagged." The possible values are: • **Tagged**: all frames transmitted for this VLAN will be tagged. • **Untagged**: all frames transmitted for this VLAN will be untagged.

If you make any changes to the page, click **Submit** to apply the changes to the system. To delete a VLAN, select the VLAN from the **VLAN ID and Name** field, and click **Delete**. You cannot delete the default VLAN.

3.2.2 VLAN Status

Use the VLAN Status page to view information about the VLANs configured on your system.

To access the VLAN Status page, click **Switching > VLAN > Status** in the navigation tree.

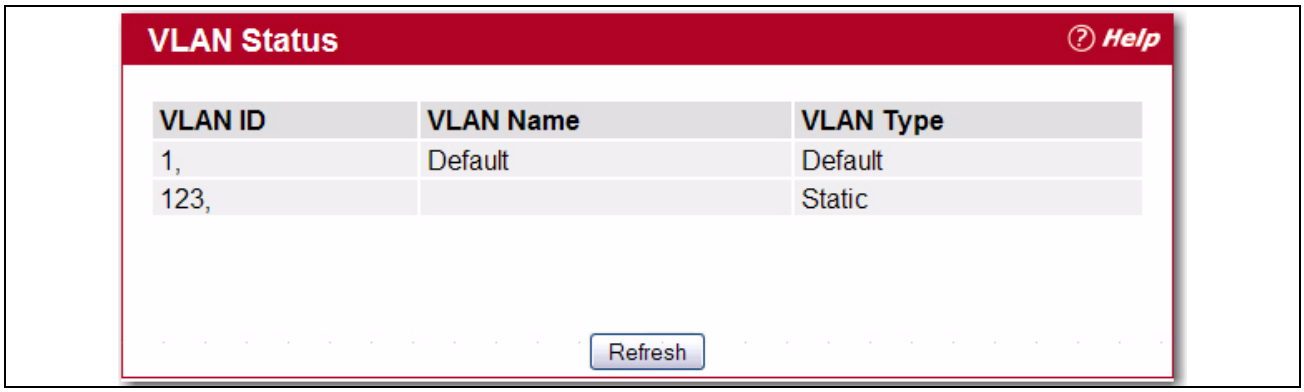

Figure 3-13: VLAN Status

Table 3-13: VLAN Status Fields

Field	Description
VLAN ID	The VLAN Identifier (VID) of the VLAN. The range of the VLAN ID is 1 to 3965.
VLAN Name	The name of the VLAN. VLAN ID 1 is always named Default.
VLAN Type	The VLAN type, which can be one of the following: • **Default**: (VLAN ID = 1) -- always present • **Static**: A VLAN you have configured • **Dynamic**: A VLAN created by GVRP registration that you have not converted to static, and that GVRP may therefore remove

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

3.2.3 VLAN Port Configuration

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

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

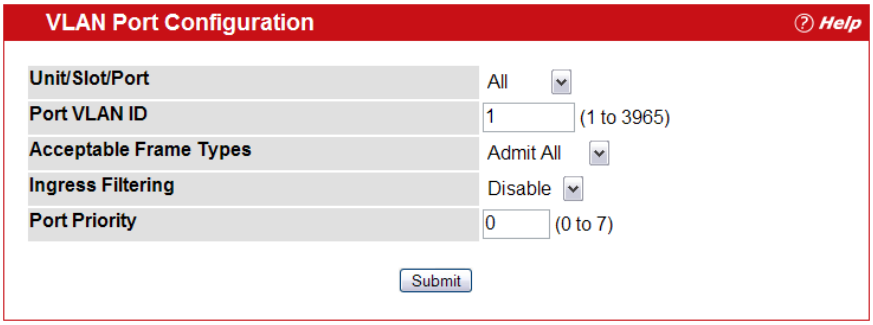

Figure 3-14: VLAN Port Configuration

Table 3-14: VLAN Port Configuration Fields

Field	Description
Slot/Port	Select the physical interface for which you want to display or configure data. Select **All** to set the parameters for all ports to same values.
Port VLAN ID	Specify the VLAN ID you want assigned to untagged or priority tagged frames received on this port. The factory default is 1.
Acceptable Frame Types	Specify how you want the port to handle untagged and priority tagged frames. Whichever you select, VLAN tagged frames will be forwarded in accordance with the IEEE 802.1Q VLAN standard. The factory default is Admit All. • **VLAN Only**: The port will discard any untagged or priority tagged frames it receives. • **Admit All**: Untagged and priority tagged frames received on the port will be accepted and assigned the value of the Port VLAN ID for this port.
Ingress Filtering	Specify how you want the port to handle tagged frames: • **Enable**: A tagged frame will be discarded if this port is not a member of the VLAN identified by the VLAN ID in the tag. • **Disable**: All tagged frames will be accepted. The factory default is disable.
Port Priority	Specify the default 802.1p priority assigned to untagged packets arriving at the port.

If you change any information on the page, click **Submit** to apply the changes to the system.

3.2.4 VLAN Port Summary

Use the VLAN Port Summary page to view VLAN configuration information for all the ports on the system.

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

VLAN Port Summary

Help

List of all Ports on the Switch

Interface	Port VLAN ID Configured	Acceptable Frame Types	Ingress Filtering Configured	Port Priority
0/1	1	Admit All	Disable	0
0/2	1	Admit All	Disable	0
0/3	1	Admit All	Disable	0
0/4	1	Admit All	Disable	0
0/5	1	Admit All	Disable	0
0/6	1	Admit All	Disable	0
0/7	1	Admit All	Disable	0
0/8	1	Admit All	Disable	0
0/9	1	Admit All	Disable	0
0/10	1	Admit All	Disable	0
0/11	1	Admit All	Disable	0
0/12	1	Admit All	Disable	0
0/13	1	Admit All	Disable	0
0/14	1	Admit All	Disable	0
0/15	1	Admit All	Disable	0

Figure 3-15: VLAN Port Summary

Table 3-15: VLAN Port Summary Fields

Field	Description
Slot/Port	Identifies the physical interface associated with the rest of the data in the row.
Port VLAN ID Configured	Identifies the VLAN ID assigned to untagged or priority-tagged frames received on this port. The factory default is 1.
Port VLAN ID Current	Displays the actual VLAN ID in use for the port. If the port was acquired by another module, the actual value may differ from the configured VLAN ID. For example, if the port is a member of a port channel and the port channel has a different port VLAN ID setting than the configured value, the two may differ.
Acceptable Frame Types	Indicates how the port handles untagged and priority tagged frames. • **VLAN Only**: The port discards any untagged or priority tagged frames it receives. • **Admit All**: Untagged and priority tagged frames received on the port are accepted and assigned the value of the Port VLAN ID for this port.
Ingress Filtering	Shows how the port handles tagged frames. • **Enable**: A tagged frame is discarded if this port is not a member of the VLAN identified by the VLAN ID in the tag. • **Disable**: All tagged frames are accepted, which is the factory default.
Port Priority	Identifies the default 802.1p priority assigned to untagged packets arriving at the port.

Click **Refresh** to reload the page and view the most current information.

3.2.5 Reset VLAN Configuration

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

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

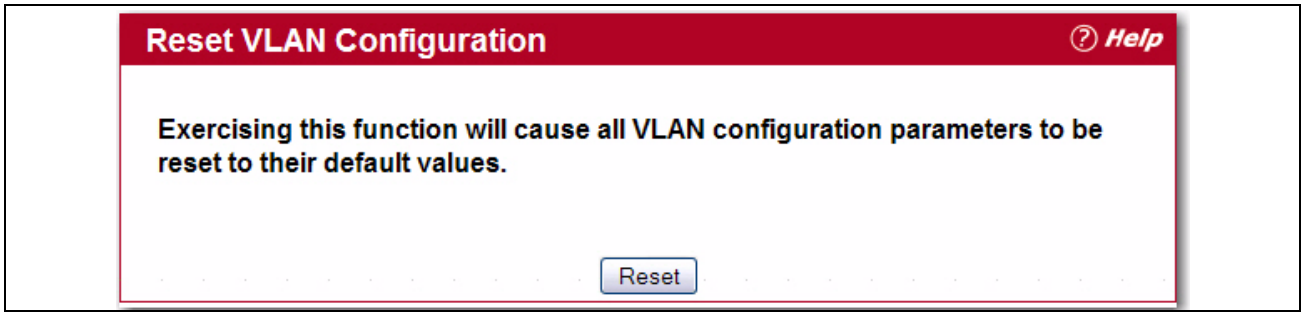

Figure 3-16: Reset VLAN Configuration

When you click **Reset**, the screen refreshes, and you are asked to confirm the reset. Click **Reset** again to restore all default VLAN settings for the ports on the system.

3.2.6 VLAN Internal Usage Configuration

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

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

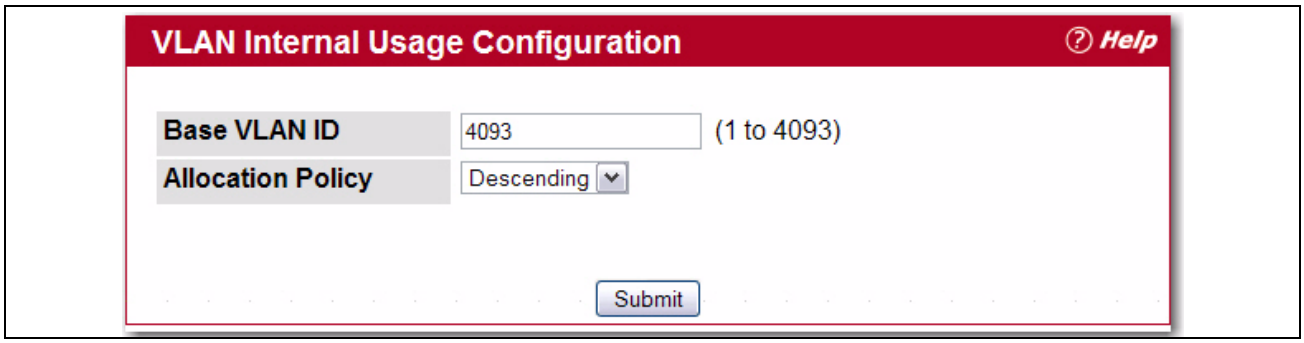

Figure 3-17: VLAN Internal Usage Configuration

Table 3-16: VLAN Internal Usage Configuration Fields

Field	Description
Base VLAN ID	The Base VLAN ID for Internal allocation of VLANs to the routing interface.
Allocation Policy	Allocation Policy for VLAN ID in ascending or descending order.

If you change any information on the page, click **Submit** to apply the changes to the system.

3.3 Double VLAN (DVLAN) Tunneling

DVLAN Tunneling allows the use of a second tag on network traffic. The additional tag helps differentiate between customers in the Metropolitan Area Networks (MAN) while preserving individual customer's VLAN identification when they enter their own 802.1Q domain.

With the introduction of this second tag, you do not need to divide the 4k VLAN ID space to send traffic on an Ethernet-based MAN.

With DVLAN Tunneling enabled, every frame that is transmitted from an interface has a new VLAN tag (S-tag) attached while every packet that is received from an interface has a VLAN tag (S-tag) removed (if one or more tags are present).

DVLAN also supports up to 4 Tag Protocol Identifier (TPID) values per switch and the ability to map these values to ports. This allows you to configure the same or different TPIDs for different ports.

Use the DVLAN Tunneling page to configure DVLAN frame tagging on one or more ports.

The DVLAN folder contains links to the following features:

- DVLAN Config
- DVLAN Summary
- DVLAN Interface Config
- DVLAN Interface Summary

3.3.1 DVLAN Config

The DVLAN **Config** page allows you to configure the TPID with an associated Global Ethertype for all ports on the system.

To access the DVLAN **Config** page, click **Switching > DVLAN > Config** in the navigation tree.

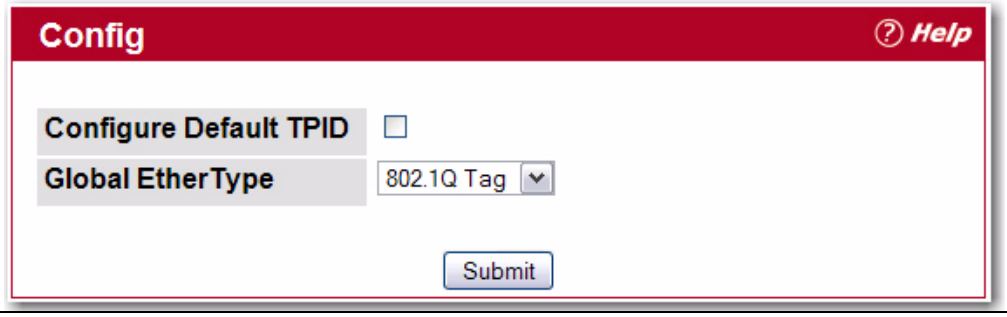

Figure 3-18: DVLAN Config

Table 3-17: DVLAN Config Fields

Field	Description
Configure Default TPID	Select to configure the selected Global Default EtherType.
Global EtherType	Specifies one of the following global EtherType options: • 802.1Q Tag - Commonly used tag representing 0x8100 • vMAN Tag - Commonly used tag representing 0x88A8 • Custom Tag - Configure the tag for EtherType by providing a custom value in any range from 0 to 65535. The two-byte hex EtherType is used as the first 16 bits of the DVLAN tag.

If you make any changes to the page, click **Submit** to apply the changes to the system.

3.3.2 DVLAN Summary

The DVLAN **Summary** page allows you to view the Global and Default TPIDs configured for all ports on the system.

To access the DVLAN **Summary** page, click **Switching > DVLAN > Summary** in the navigation tree.

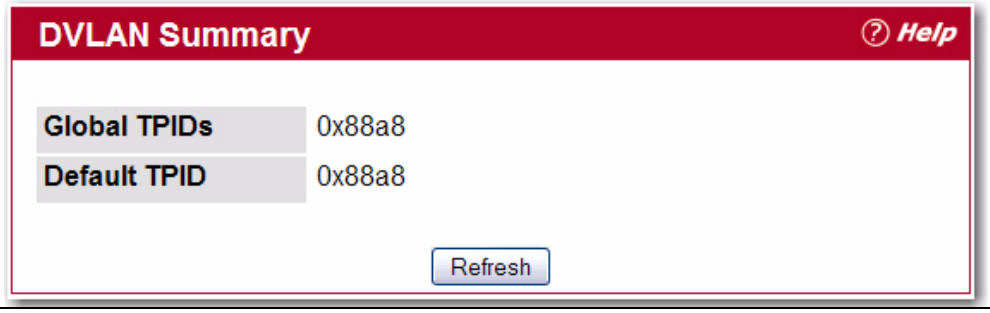

Figure 3-19: DVLAN Summary

Table 3-18: DVLAN Summary Fields

Field	Description
Global TPIDs	Displays the Global tag protocol identifiers configured.
Default TPID	Displays the Default tag protocol identifiers configured.

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

3.3.3 DVLAN Interface Config

The DVLAN **Interface Config** page allows you to view and configure the DVLAN interface configuration status for all ports on the system.

To access the DVLAN **Interface Config** page, click **Switching > DVLAN > Interface Config** in the navigation tree.

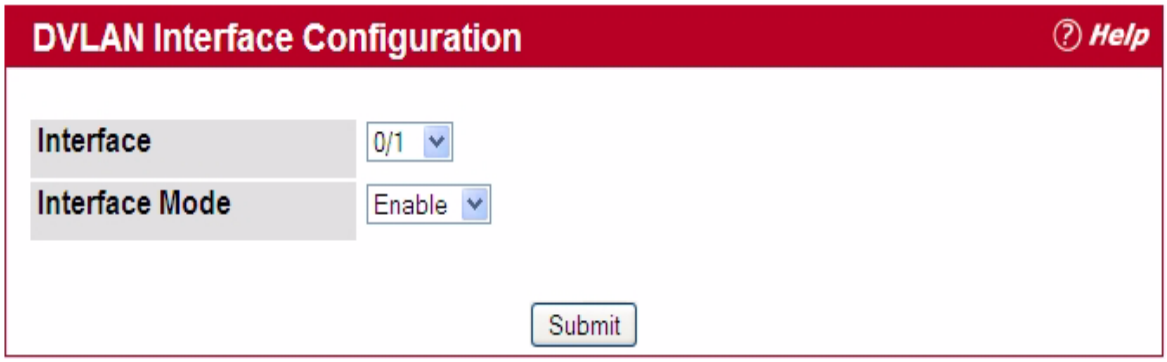

Figure 3-20: DVLAN Interface Config

Table 3-19: DVLAN Interface Config Fields

Field	Description
Slot/Port	Select the physical interface for which you want to display or configure data. Select **All** to set the parameters for all ports to same values.
Interface Mode	This specifies the administrative mode for DVLAN Tagging: • **Enable**: DVLAN Tagging is enabled for the specified port (or All ports). • **Disable**: DVLAN Tagging is disabled for the specified port (or All ports), which is the default value.

Click **Refresh** to redisplay the most current information from the router.

3.3.4 DVLAN Interface Summary

The DVLAN **Interface Summary** page displays the DVLAN interface configuration status for all ports on the system.

To access the DVLAN **Interface Summary** page, click **Switching > DVLAN > Interface Summary** in the navigation tree.

DVLAN Interface Summary Help

Interface	Interface Mode	Interface EtherType
0/1	Enable	0x8100
0/2	Disable	0x8100
0/3	Disable	0x8100
0/4	Disable	0x8100
0/5	Disable	0x8100
0/6	Disable	0x8100
0/7	Disable	0x8100
0/8	Disable	0x8100
0/9	Disable	0x8100
0/10	Disable	0x8100
0/11	Disable	0x8100
0/12	Disable	0x8100
0/13	Disable	0x8100
0/14	Disable	0x8100
0/15	Disable	0x8100

Figure 3-21: DVLAN Interface Summary

See 3.3.3 DVLAN Interface Config for a description of these fields.

Click **Refresh** to redisplay the most current information from the router.

3.4 Configuring Protected Ports

The Protected Ports feature assists in Layer 2 security. Ports that are configured to be protected cannot forward traffic to other protected ports in the same group, regardless of having the same VLAN membership. However, protected ports can forward traffic to ports which are unprotected as well as ports in other protected groups. Unprotected ports can forward traffic to both protected and unprotected ports.

3.4.1 Protected Port Configuration

Use the Protected Ports Configuration page to create up to three protected port groups and to assign physical ports to a group.

To display the Protected Port Configuration page, click **Switching > Protected Ports > Configuration** in the navigation tree.

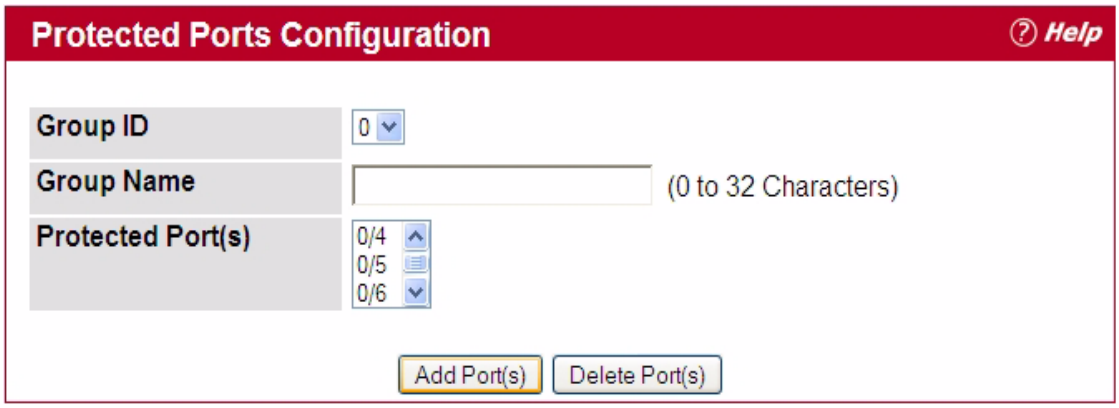

Figure 3-22: Protected Port Configuration

Table 3-20: Protected Port Configuration Fields

Field	Description
Group ID	The protected ports can be combined into a logical group. Traffic can flow between protected ports belonging to different groups, but not within the same group. The selection box lists all the possible protected port Group IDs supported for the current platform. The valid range is platform-dependent.
Group Name	Assign an optional name to associate with the protected ports group. The name is for identification purposes and can be up to 32 characters long, including blanks. The default is blank.
Protected Port(s)	Specifies the Slot and Port for which port parameters are defined.

3.4.1.1 Assigning Ports to a Group

1. Select a group ID from the **Group ID** field.
2. From the **Protected Port(s)** field, click one port to add a single port to the group, or hold the CTRL key and click multiple ports to add more than one port to the group.
3. Click **Submit** to apply the changes to the system.

3.4.2 Protected Ports Summary

Use the Protected Ports Summary page to view information about protected port groups and their included ports.

To view the Protected Ports Summary page, click **Switching > Protected Ports > Summary** in the navigation tree.

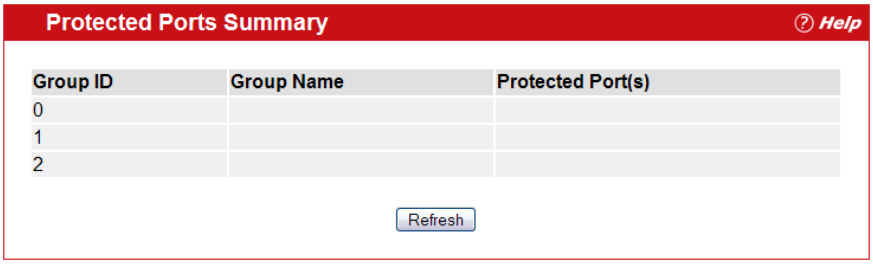

Figure 3-23: Protected Ports Summary

Table 3-21: Protected Ports Summary Fields

Field	Description
Group ID	Identifies the protected ports group as either Group 0, 1, or 2.
Group Name	Identifies the protected ports group with a user-defined string.
Protected Port(s)	Shows the Slot and Port that are members of the protected ports group.

Click **Refresh** to reload the page and display the most current information.

3.5 Managing Protocol-Based VLANs

In a protocol-based VLAN, traffic is bridged through specified ports based on the protocol associated with the VLAN. User-defined packet filters determine whether a particular packet belongs to a particular VLAN. Protocol-based VLANs are most often used in situations where network segments contain hosts running multiple protocols.

You can use a protocol-based VLAN to define filtering criteria for untagged packets. By default, if you do not configure any port-based (IEEE 802.1Q) or protocol-based VLANs, untagged packets are assigned to VLAN 1. You can override this behavior by defining either port-based VLANs, protocol-based VLANs, or both. Tagged packets are always handled according to the IEEE 802.1Q standard and are not included in protocol-based VLANs.

If you assign a port to a protocol-based VLAN for a specific protocol, untagged frames received on that port for that protocol will be assigned the protocol-based VLAN ID. Untagged frames received on the port for other protocols will be assigned the Port VLAN ID (PVID), which is either the default PVID (1) or a PVID you have specifically assigned to the port using the Port VLAN Configuration screen.

3.5.1 Configuration

Use the Protocol-based VLAN Configuration page to configure which protocols go to which VLANs, and enable certain ports to use these settings.

You define a protocol-based VLAN by creating a group. Each group has a one-to-one relationship with a VLAN ID, can include one or more protocol definitions (the range is platform-dependent), and can include multiple ports.

To display the Protocol-Based VLAN Configuration page, click **Switching > VLAN > Protocol-based VLAN > Configuration** in the navigation tree.

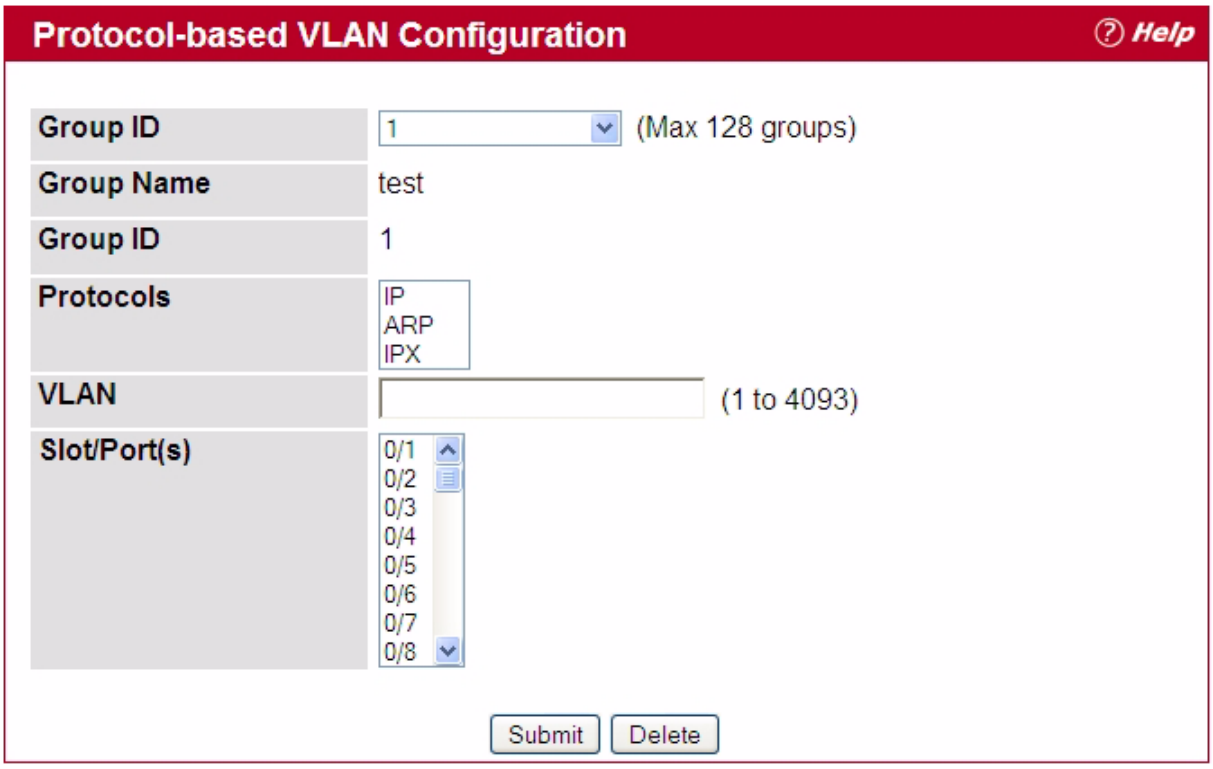

Figure 3-24: Protocol Group

Table 3-22: Protocol Group Fields

Field	Description
Group	Use the drop-down menu to create or modify a protocol group. You can create up to 128 groups.
Group Name	When creating a group, enter a name to associate with protocol group ID. You can modify the name of an existing group. You can enter up to 16 characters.
Group ID	Shows the number that identifies the group you create. Group IDs are automatically assigned when you create a group.
Protocols	Select one or more protocols to associate with this group. CTRL + click to select multiple protocols. • **IP**: IP is a network layer protocol that provides a connectionless service for the delivery of data. • **ARP**: Address Resolution Protocol (ARP) is a low-level protocol that dynamically maps network layer addresses to physical medium access control (MAC) addresses • **IPX**: The Internetwork Packet Exchange (IPX) is a connectionless datagram Network-layer protocol that forwards data over a network.
VLAN ID	Specifies the VLAN ID associated with this group. The range is 1-3965.
Slot/Port	Selects the interface(s) to add or remove from this group. CTRL + click to select multiple protocols.

- To create or modify a protocol-based VLAN group, edit the fields, and click **Submit**.
- To delete an existing protocol-based VLAN group, select the group from the **Group ID** field, and click **Delete Group**.

3.5.2 Protocol-Based VLAN Summary

Use the Protocol-based VLAN Summary page to view information about protocol-based VLAN groups configured on the system.

To access the Protocol-based VLAN Summary page, click **Switching > Protocol-based VLAN > Summary** in the navigation tree.

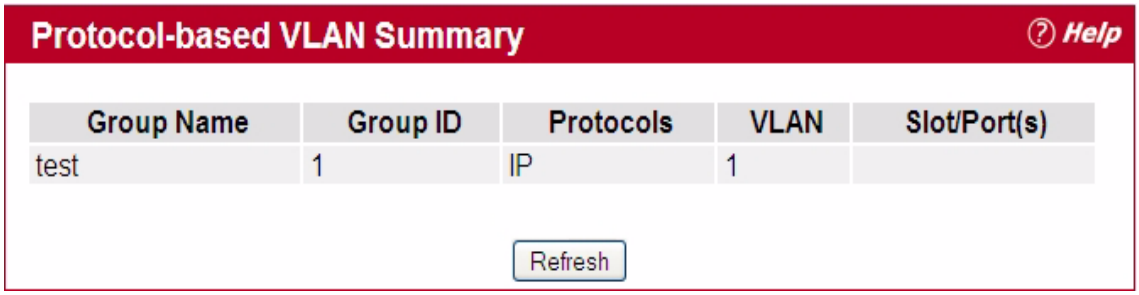

Figure 3-25: Protocol-based VLAN Summary

Table 3-23: Protocol-based VLAN Summary Fields

Field	Description
Group Name	Shows the user-defined name associated with protocol group.
Group ID	Shows the number that identifies the group you create. Group IDs are automatically assigned when you create a group.
Protocols	Shows the protocols to associate with this group, which can be one or more of the following: • **IP**: IP is a network layer protocol that provides a connectionless service for the delivery of data. • **ARP**: Address Resolution Protocol (ARP) is a low-level protocol that dynamically maps network layer addresses to physical medium access control (MAC) addresses • **IPX**: The Internetwork Packet Exchange (IPX) is a connectionless datagram Network-layer protocol that forwards data over a network.
VLAN	Specifies the VLAN ID associated with this group.
Slot/Port	Shows the interfaces participating in this group.

Click **Refresh** to reload the page and display the most current information.

3.6 Managing IP Subnet-Based VLANs

If a packet is untagged or priority- tagged, the device associates the packet with any matching IP subnet classification. If no IP subnet classification can be made, the packet is subjected to the normal VLAN classification rules of the device. An IP subnet-to-VLAN mapping is defined by configuring an entry in the IP subnet-to-VLAN table. An entry is specified by a source IP address, network mask, and the desired VLAN ID. The IP subnet-to-VLAN configurations are shared across all ports of the switch.

3.6.1 Configuration

Use the IP Subnet-based VLAN Configuration page to assign an IP Subnet to a VLAN.

To display the IP Subnet-based VLAN Configuration page, click **Switching > IP Subnet-based VLAN > Configuration** in the navigation menu.

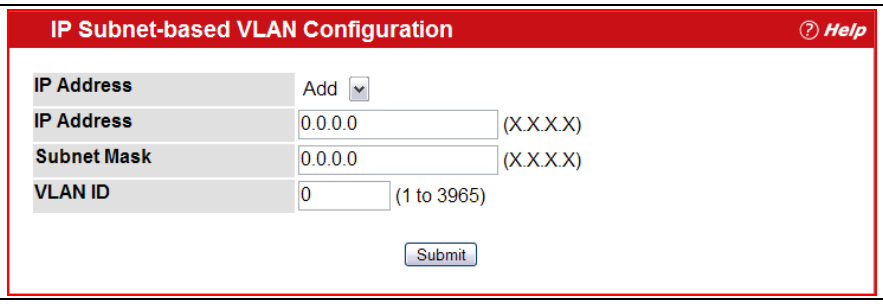

Figure 3-26: IP Subnet-based VLAN Configuration

Table 3-24: IP Subnet-based VLAN Configuration Fields

Field	Description
IP Address	Select the IP address of the IP-to-VLAN binding to view or delete, or select **Add** to create a new binding.
IP Address	Specifies packet source IP address. This field is configurable only when you create a new IP Subnet-based VLAN. Enter the IP address in dotted decimal notation.
Subnet Mask	Specifies packet source IP subnet mask address.This field is configurable only when you create a new IP Subnet-based VLAN. Enter the subnet mask in dotted decimal notation.
VLAN ID	Specifies the VLAN to which the IP address is assigned. The valid range is 1-3965.

- If you make any changes on this page, click **Submit** to apply the changes to the system.
- To delete an existing binding, select the source IP address from the **IP Address** drop-down menu, and click **Delete**.

3.6.2 Summary

Use the IP Subnet-based VLAN Summary page to view information about IP subnet to VLAN mappings configured on your system. If no mappings are configured, the screen displays a “No IP Subnet-based VLAN Configured” message.

To access the IP Subnet-based VLAN Summary page, click **Switching > IP Subnet-based VLAN Summary** in the navigation tree.

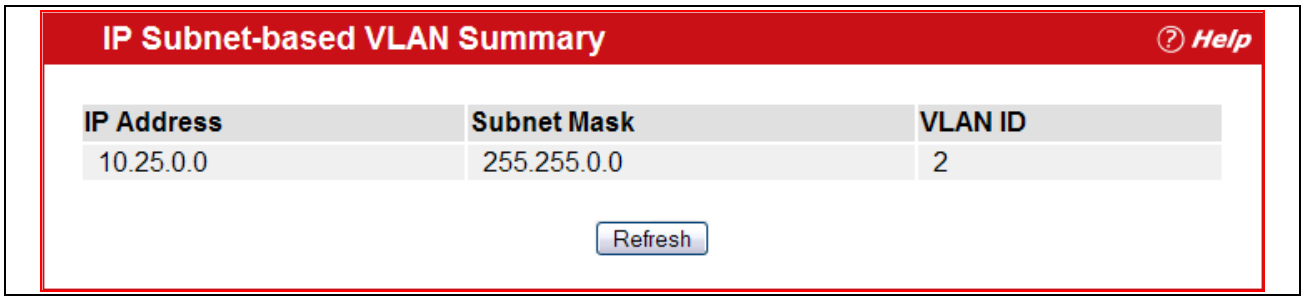

Figure 3-27: IP Subnet-based VLAN Summary

Table 3-25: IP Subnet-based VLAN Summary Fields

Field	Description
IP Address	Shows the packet source IP address.
Subnet Mask	Shows packet source IP subnet mask address.
VLAN ID	Shows the VLAN to which the IP address is assigned.

Click **Refresh** to reload the page and display the most current information.

3.7 Managing MAC-Based VLANs

3.7.1 MAC-based VLAN Configuration

If a packet is untagged or priority tagged, the device shall associate it with the VLAN which corresponds to the source MAC address in its MAC-based VLAN tables. If there is no matching entry in the table, the packet is subject to normal VLAN classification rules of the device.

Use the MAC-based VLAN Configuration page to map a MAC entry to the VLAN table. After the source MAC address and the VLAN ID are specified, the MAC-to-VLAN configurations are shared across all ports of the switch.

To display the MAC-based VLAN Configuration page, click **Switching > MAC-based VLAN > Configuration** in the navigation menu.

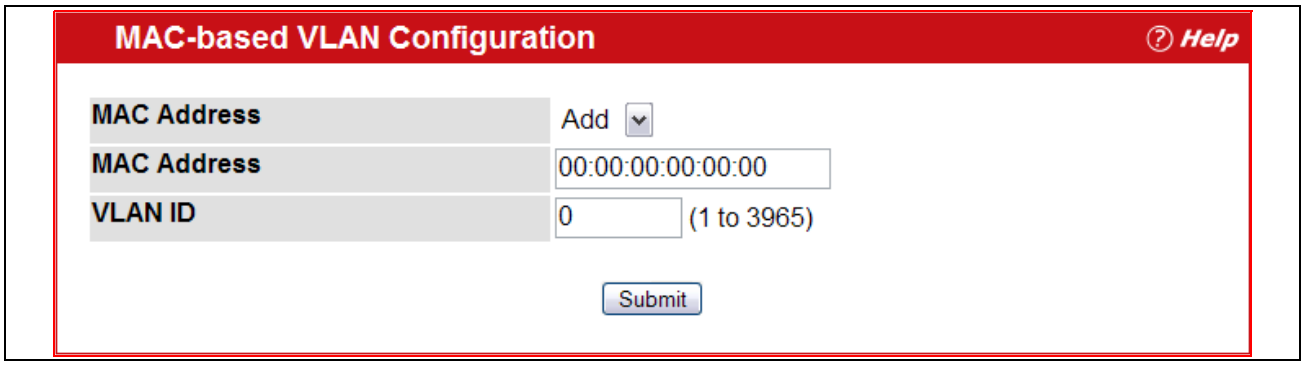

Figure 3-28: MAC-based VLAN Configuration

Table 3-26: MAC-based VLAN Configuration Fields

Field	Description
MAC Address	Specifies the source MAC address to map to a VLAN.
VLAN ID	Specifies the VLAN to which the source MAC address is to be bound.

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

3.7.2 MAC-based VLAN Summary

Use the MAC-based VLAN Summary page to view information about the MAC-to-VLAN mappings configured on your system.

To display the MAC-based VLAN Summary page, click **Switching > MAC-based VLAN > Summary** in the navigation menu.

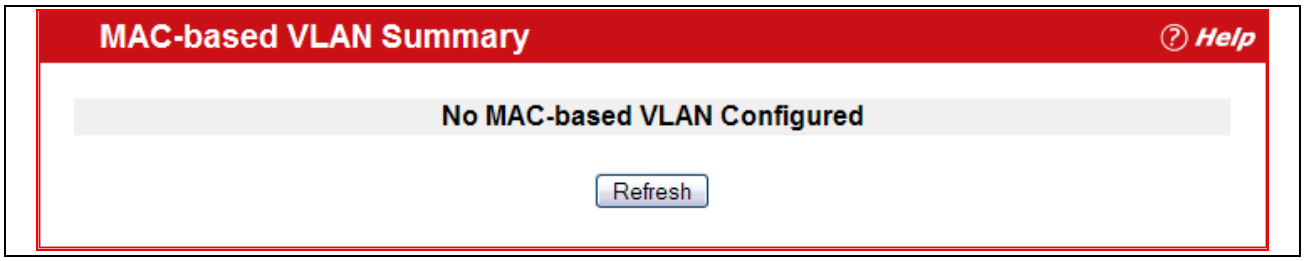

Figure 3-29: MAC-based VLAN Summary

Table 3-27: MAC-based VLAN Summary Fields

Field	Description
MAC Address	Specifies the MAC address to map to a VLAN.
VLAN ID	Specifies the VLAN to which the MAC is to be bound.

Click **Refresh** to reload the page and display the most current information.

3.8 Voice VLAN Configuration

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

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

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

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

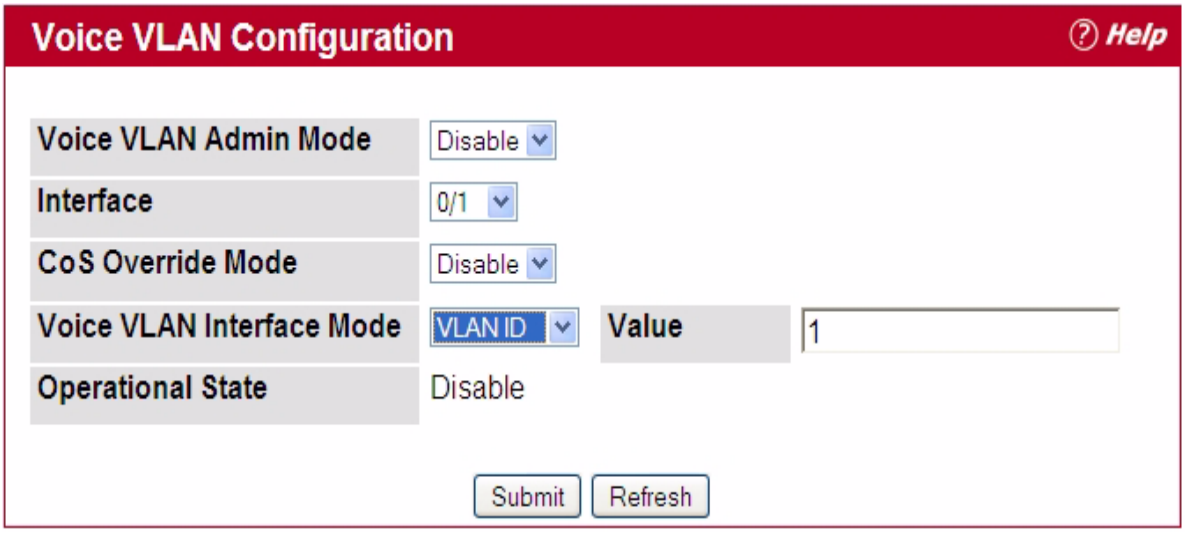

Figure 3-30: Voice VLAN Configuration

Table 3-28: Voice VLAN Configuration Fields

Field	Description
Voice VLAN Admin Mode	Click **Enable** or **Disable** to administratively turn the Voice VLAN feature on or off for all ports.
Slot/Port	Select the slot and port to configure this service on.
Voice VLAN Interface Mode	Select one of the following interface modes: • **Disable**: The voice VLAN service is disabled on this interface. Note that the Admin mode field takes precedence; i.e., if a particular interface is enabled, but the Admin Mode field is set to Disabled, the service will not be operational. • **None**: The voice VLAN service is disabled on this interface; however, unlike Disable mode, the CoS override feature is still operational on the port. • **VLAN ID**: The voice VLAN packets are uniquely identified by a number you assign. All voice traffic carries this VLAN ID to distinguish it from other data traffic which is assigned the port's default VLAN ID. However, voice traffic is not prioritized differently than other traffic. • **dot1p**: This parameter is set by the VoIP device for all voice traffic to distinguish voice data from other traffic. All other traffic is assigned the port's default VLAN ID. This feature may not be supported by all hardware configurations. • **Untagged**:
CoS Override Mode	Overrides the 802.1p class-of-service (CoS) value for all data (non-voice) packets arriving at the port. Thus any rogue client that is also connected to the voice VLAN port cannot deteriorate the voice traffic.
Operational State	Indicates whether the voice VLAN is operational.

- If you make any changes, click **Submit** to apply the change to the system.
- Click **Refresh** to display the latest information from the router.

3.9 Creating MAC Filters

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

3.9.1 MAC Filter Configuration

Use the MAC Filter Configuration page to associate a MAC address with a VLAN and one or more source and/or destination ports

To access the MAC Filter Configuration page, click **Switching > Filters > Configuration** in the navigation tree.

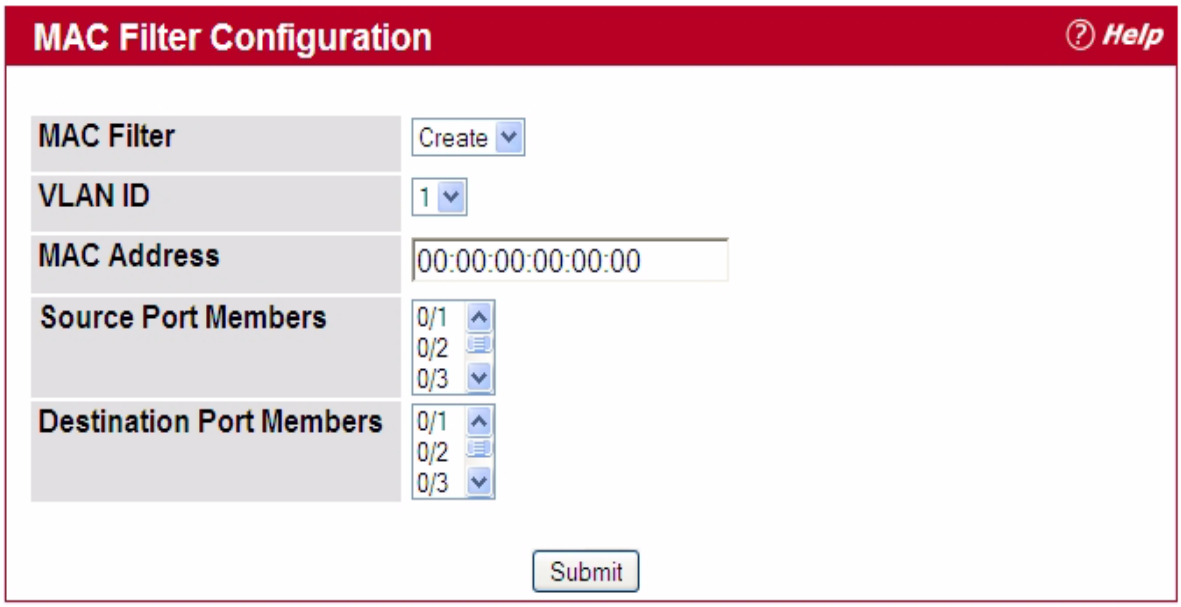

Figure 3-31: MAC Filter Configuration

Table 3-29: MAC Filter Configuration Fields

Field	Description
MAC Filter	If no MAC filters are configured on the system, **Create Filter** is the only item in the drop-down menu. If one or more MAC filters exist, the list also contains the MAC address and associated VLAN ID of a configured filter.
MAC Address	The MAC address of the filter in the format 00:01:1A:B2:53:4D. You can only change this field when you have selected the "Create Filter" option. **Note:** You cannot define filters for the following MAC addresses: • 00:00:00:00:00:00 • 01:80:C2:00:00:00 to 01:80:C2:00:00:0F • 01:80:C2:00:00:20 to 01:80:C2:00:00:21 • FF:FF:FF:FF:FF:FF
VLAN ID	The VLAN ID used with the MAC address to fully identify packets you want filtered. You can only change this field when you have selected the "Create Filter" option.
Source Port Mask	Select the ports you want included in the inbound filter. If a packet with the MAC address and VLAN ID you selected is received on a port that is not in the list, it will be dropped.
Destination Port Mask	Select the ports you want to include in the outbound filter. A packet, once admitted, is sent to all the ports in this list.

3.9.1.1 Adding MAC Filters

1. To add a MAC filter, select **Create Filter** from the **MAC Filter** drop-down menu.
2. Enter a valid MAC address and select a VLAN ID from the drop-down menu.

   The VLAN ID drop-down menu only lists VLANs currently configured on the system.
3. Select one or more ports to include in the filter. Use **CTRL** + click to select multiple ports.

4. Click **Submit** to apply the changes to the system.

3.9.1.2 Modifying MAC Filters

To change the port mask(s) for an existing filter, select the entry from the **MAC Filter** field, and click (or CTRL + click) the port(s) to include in the filter. Only those ports that are highlighted when you click **Submit** are included in the filter.

To change the MAC address or VLAN associated with a filter, you must delete and re-create the filter.

3.9.1.3 Deleting MAC Filters

To delete a filter, select it from the **MAC Filter** drop-down menu and click **Delete**. To delete all configured filters from the forwarding database, click **Delete All**.

3.9.2 MAC Filter Summary

Use the MAC Filter Summary page to associate a MAC address with a VLAN and one or more source ports.

To access the MAC Filter Summary page, click **Switching > Filters > Summary** in the navigation tree.

Figure 3-32: MAC Filter Summary

Table 3-30: MAC Filter Summary Fields

Field	Description
MAC Address	Shows the MAC address of the filter.
VLAN ID	Shows the VLAN ID used with the MAC address to fully identify packets you want filtered.
Source Port Members	Lists the ports included in the inbound filter. If a packet with the MAC address and VLAN ID displayed is received on a port that is not in the list, it will be dropped.
Destination Port Members	Lists the ports included in the outbound filter. A packet, once admitted, is sent to all ports in the list.

3.10 Configuring GARP

Generic Attribute Registration Protocol (GARP) is a general-purpose protocol that registers any network connectivity or membership-style information. GARP defines a set of switches interested in a given network attribute, such as VLAN or multicast address.

The GARP VLAN Registration Protocol (GVRP) provides a mechanism that allows networking switches to dynamically register (and de-register) VLAN membership information with the networking devices attached to the same segment, and for that information to be disseminated across all networking switches in the bridged LAN that support GVRP.

With the GARP Multicast Registration Protocol (GMRP), networking devices can dynamically register and de-register group membership information with the networking devices attached to the same segment. GMRP enables the group membership information to be disseminated across all networking devices in the bridged LAN that support Extended Filtering Services.

The operation of GVRP and GMRP relies upon the services provided by GARP.

3.10.1 GARP Status

Use the GARP Status page to view GARP settings for the system and for each interface.

To access the GARP Status page, click **Switching > GARP > Status** in the navigation tree.

GARP Status Help

Switch GVRP	Disable
Switch GMRP	Disable

Interface	Port GVRP Mode	Port GMRP Mode	Join Timer (centisecs)	Leave Timer (centisecs)	Leave All Timer (centisecs)
0/1	Disable	Disable	20	60	1000
0/2	Disable	Disable	20	60	1000
0/3	Disable	Disable	20	60	1000
0/4	Disable	Disable	20	60	1000
0/5	Disable	Disable	20	60	1000
0/6	Disable	Disable	20	60	1000
0/7	Disable	Disable	20	60	1000
0/8	Disable	Disable	20	60	1000
0/9	Disable	Disable	20	60	1000
0/10	Disable	Disable	20	60	1000
0/11	Disable	Disable	20	60	1000
0/12	Disable	Disable	20	60	1000
0/13	Disable	Disable	20	60	1000
0/14	Disable	Disable	20	60	1000
0/15	Disable	Disable	20	60	1000

Figure 3-33: GARP Status

The GARP Status page contains the following fields:

Table 3-31: GARP Status Fields

Field	Description
Switch GVRP	Shows whether the switch GVRP protocol is enabled or disabled.
Switch GMRP	Shows whether the switch GMRP protocol is enabled or disabled.
Slot/Port	Identifies the system interface.
Port GVRP Mode	Shows the GARP VLAN Registration Protocol administrative mode for the port. If the mode is Disabled, the protocol will not be active and the Join Time, Leave Time and Leave All Time will have no effect.

Table 3-31: GARP Status Fields

Field	Description
Port GMRP Mode	Shows the GARP Multicast Registration Protocol administrative mode for the port. If the mode is Disabled, the protocol will not be active, and Join Time, Leave Time and Leave All Time have no effect.
Join Timer (centisecs)	Shows the time between the transmission of GARP PDUs registering (or re-registering) membership for a VLAN or multicast group in centiseconds.
Leave Timer (centisecs)	Displays time lapse, in centiseconds, that the switch waits before leaving its GARP state. Leave time is activated by a Leave All Time message sent/received, and cancelled by the Join message received.This allows time for another station to assert registration for the same attribute in order to maintain uninterrupted service.
Leave All Timer (centisecs)	Displays time lapse, in centiseconds, that all switches wait before leaving the GARP state. The leave all time must be greater than the leave time. The Leave All Time controls how frequently LeaveAll PDUs are generated. A LeaveAll PDU indicates that all registrations will shortly be deregistered. Participants will need to rejoin in order to maintain registration.

3.10.2 GARP Switch Configuration

Use the GARP Switch Configuration page to configure GARP settings for the system.

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

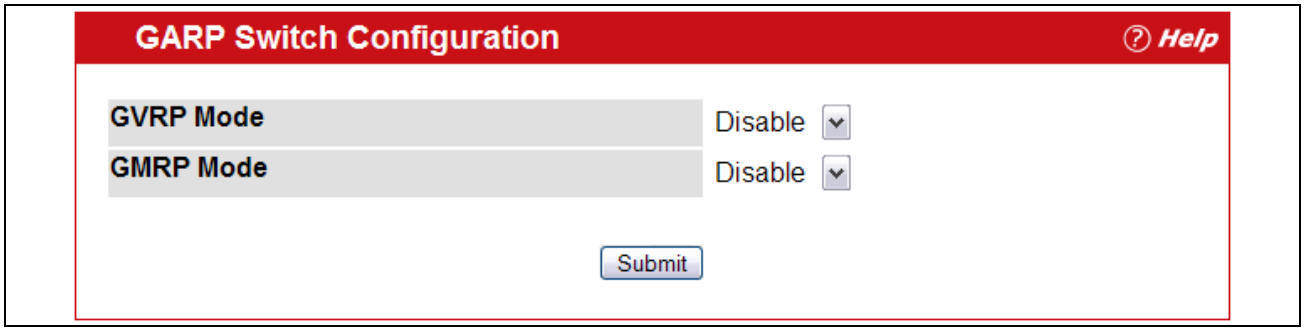

Figure 3-34: GARP Switch Configuration

Table 3-32: GARP Switch Configuration Fields

Field	Description
Switch GVRP Mode	Shows the GARP VLAN Registration Protocol administrative mode for the switch. The switch GVRP mode must be enabled for the ports to function in GARP protocols, even if GVRP is enabled on a port.
Switch GMRP Mode	Shows the GARP Multicast Registration Protocol administrative mode for the switch. The switch GMRP mode must be enabled for the ports to function in GARP protocols, even if GMRP is enabled on a port.

If you make any changes to the page, click **Submit** to apply the changes to the system.

3.10.3 GARP Port Configuration

Use the GARP Port Configuration page to configure GARP settings for a specific interface.

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

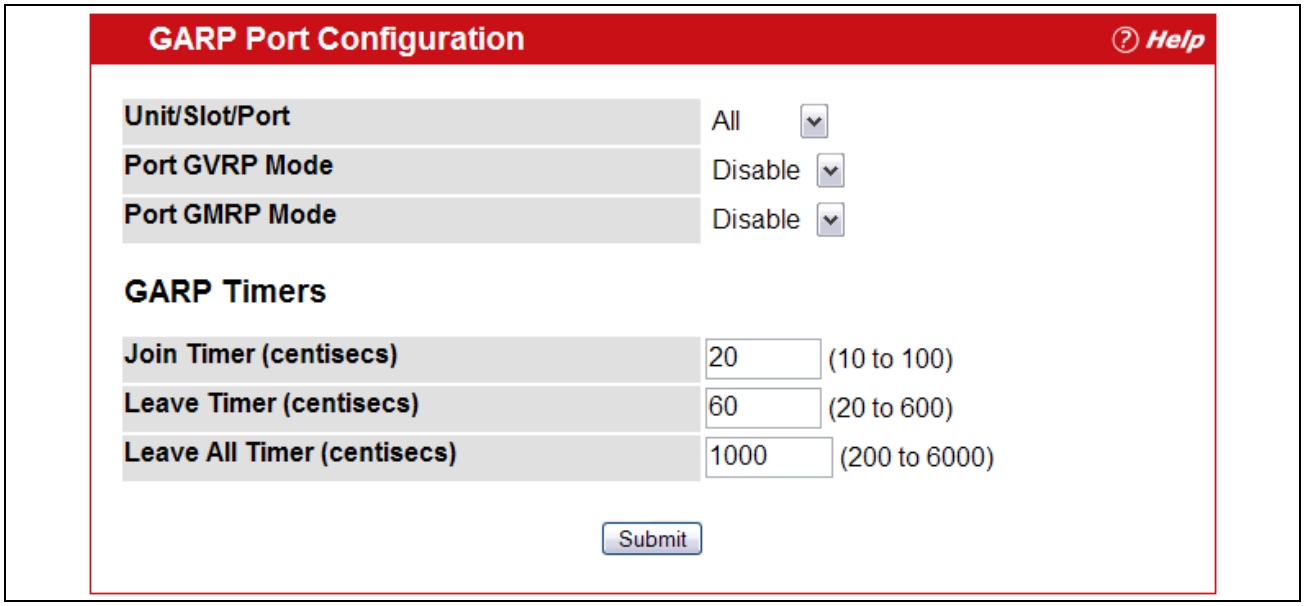

Figure 3-35: GARP Port Configuration

Table 3-33: GARP Port Configuration Fields

Field	Description
Slot/Port	Specifies interface on which to configure the GARP settings. If you select All from the drop-down menu, the settings on the page affect all interfaces.
Port GVRP Mode	Choose the GARP VLAN Registration Protocol administrative mode for the port by selecting enable or disable from the pulldown menu.If you select disable, the protocol will not be active and the Join Time, Leave Time and Leave All Time will have no effect. The factory default is disable.
Port GMRP Mode	Choose the GARP Multicast Registration Protocol administrative mode for the port by selecting enable or disable from the pulldown menu. If you select disable, the protocol will not be active, and Join Time, Leave Time and Leave All Time have no effect. The factory default is disable.

Table 3-33: GARP Port Configuration Fields (Continued)

Field	Description
	GARP Timers
GARP Join Timer (centisecs)	Specify the time between the transmission of GARP PDUs registering (or re-registering) membership for a VLAN or multicast group in centiseconds. Enter a number between 10 and 100 (0.1 to 1.0 seconds). The factory default is 20 centiseconds (0.2 seconds). An instance of this timer exists for each GARP participant for each port.
GARP Leave Timer (centisecs)	Displays time lapse, in centiseconds, that the switch waits before leaving its GARP state. Leave time is activated by a Leave All Time message sent/received, and cancelled by the Join message received.This allows time for another station to assert registration for the same attribute in order to maintain uninterrupted service. Enter a number between 20 and 600 (0.2 to 6.0 seconds). Leave time must be greater than or equal to three times the join time. The factory default is 60 centiseconds (0.6 seconds). An instance of this timer exists for each GARP participant for each port.
GARP Leave All Timer (centisecs)	Displays time lapse, in centiseconds, that all switches wait before leaving the GARP state. The leave all time must be greater than the leave time. The possible field value is 200-6000. The default value is 1000 centisecs. The Leave All Time controls how frequently LeaveAll PDUs are generated. A LeaveAll PDU indicates that all registrations will shortly be deregistered. Participants will need to rejoin in order to maintain registration. The Leave All Period Timer is set to a random value in the range of LeaveAllTime to 1.5*LeaveAllTime. The timer is specified in centiseconds. Enter a number between 200 and 6000 (2 to 60 seconds). The factory default is 1000 centiseconds (10 seconds). An instance of this timer exists for each GARP participant for each port.

If you make any changes to the page, click **Submit** to apply the changes to the system.

3.11 Configuring Dynamic ARP Inspection

Dynamic ARP Inspection (DAI) is a security feature that rejects invalid and malicious ARP packets. DAI prevents a class of man-in-the-middle attacks, where an unfriendly station intercepts traffic for other stations by poisoning the ARP caches of its unsuspecting neighbors. The miscreant sends ARP requests or responses mapping another station's IP address to its own MAC address.

DAI relies on DHCP snooping. DHCP snooping listens to DHCP message exchanges and builds a binding database of valid {MAC address, IP address, VLAN, and interface} tuples.

When DAI is enabled, the switch drops ARP packets whose sender MAC address and sender IP address do not match an entry in the DHCP snooping bindings database. You can optionally configure additional ARP packet validation.

3.11.1 DAI Configuration

Use the DAI Configuration page to configure global DAI settings.

To display the DAI Configuration page, click **Switching > Dynamic ARP Inspection > DAI Configuration** in the navigation tree.

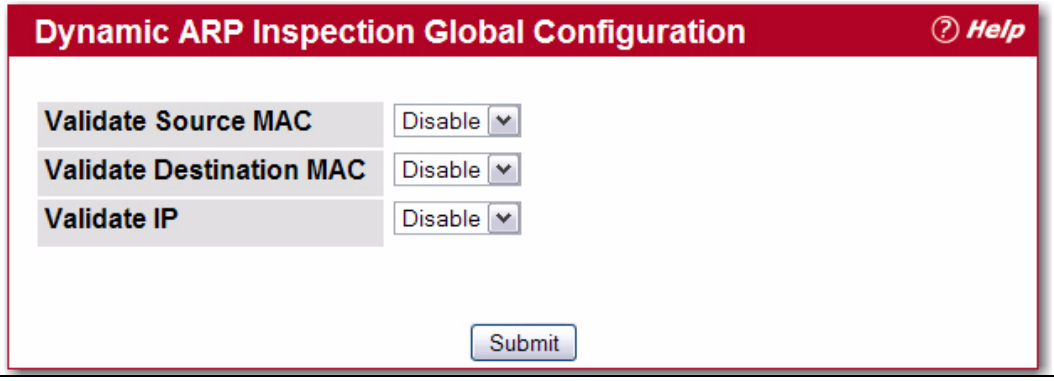

Figure 3-36: Dynamic ARP Inspection Configuration

Table 3-34: Dynamic ARP Inspection Configuration

Field	Description
Validate Source MAC	Select the DAI Source MAC Validation Mode for the switch. If you select **Enable**, Sender MAC validation for the ARP packets will be enabled. The default is **Disable**.
Validate Destination MAC	Select the DAI Destination MAC Validation Mode for the switch. If you select **Enable**, Destination MAC validation for the ARP Response packets will be enabled. The default is **Disable.**
Validate IP	Select the DAI IP Validation Mode for the switch. If you select **Enable**, IP Address validation for the ARP packets will be enabled. The default is **Disable.**

Click **Submit** to apply the new configuration and cause the change to take effect. These changes will not be retained across a power cycle unless a Save configuration is performed.

3.11.2 DAI VLAN Configuration

Use the DAI VLAN Configuration page to select the DAI-capable VLANs for which information is to be displayed or configured.

To display the DAI Configuration page, click **Switching > Dynamic ARP Inspection > DAI VLAN Configuration** in the navigation tree.

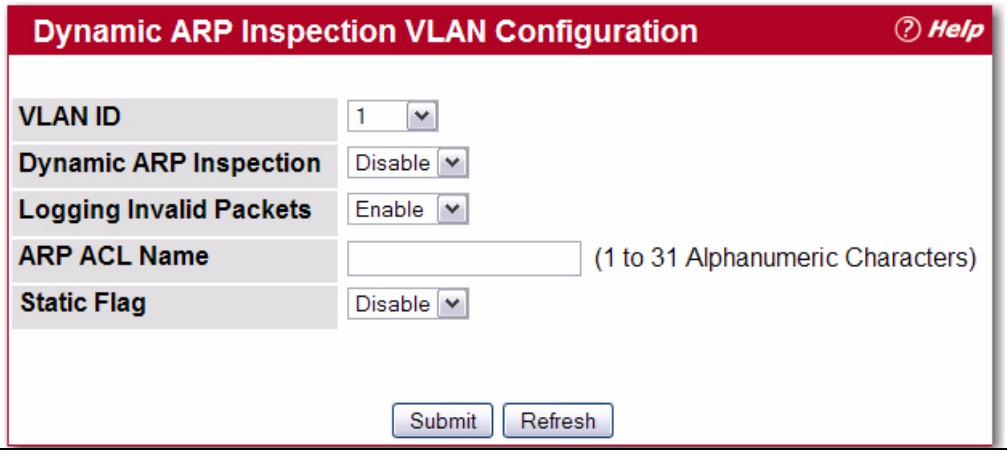

Figure 3-37: Dynamic ARP Inspection VLAN Configuration

Table 3-35: Dynamic ARP Inspection VLAN Configuration

Field	Description
VLAN ID	Select the VLAN ID for which information is to be displayed or configured.
Dynamic ARP Inspection	Select whether Dynamic ARP Inspection is **Enabled** or **Disabled** on this VLAN. The default is **Disable**.
Logging Invalid Packets	Select whether Dynamic ARP Inspection logging is **Enabled** or **Disabled** on this VLAN. The default is **Disable**.
ARP ACL Name	The name of the ARP Access List. A VLAN can be configured to use this ARP ACL containing rules as the filter for ARP packet validation. The name can contain 1-31 alphanumeric characters.
Static Flag	Use this flag to determine whether the ARP packet needs validation using the DHCP snooping database, in case the ARP ACL rules do not match. If **Enabled**, the ARP Packet will be validated by the ARP ACL Rules only. If **Disabled**, the ARP Packet needs further validation by using the DHCP Snooping entries. The default is **Disable.**

- Click **Submit** to apply the new configuration and cause the change to take effect. These changes will not be retained across a power cycle unless a Save configuration is performed.
- Click **Refresh** to refresh the page with the most current data from the switch.

3.11.3 DAI Interface Configuration

Use the DAI Interface Configuration page to select the DAI Interface for which information is to be displayed or configured.

To display the DAI Interface Configuration page, click **Switching > Dynamic ARP Inspection > DAI Interface Configuration** in the navigation tree.

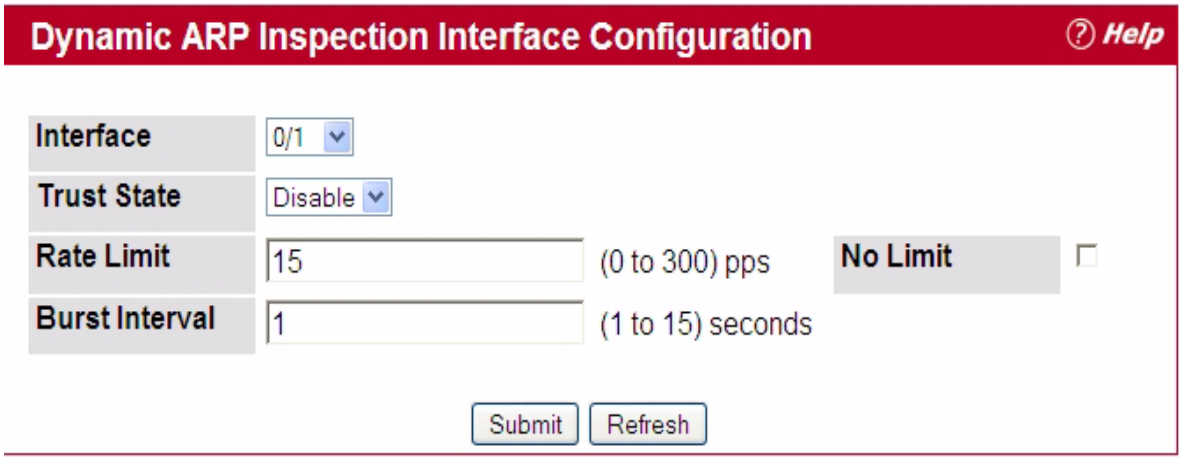

Figure 3-38: Dynamic ARP Inspection Interface Configuration

Table 3-36: Dynamic ARP Inspection Interface Configuration

Field	Description
Slot/Port	Select the physical interface for which data is to be displayed or configured.
Trust State	Indicates whether the interface is trusted for Dynamic ARP Inspection. If you select **Enable**, the interface is trusted. ARP packets coming to this interface will be forwarded without checking. If you select **Disable,** the interface is not trusted. ARP packets coming to this interface will be subjected to ARP inspection. The default is **Disable**.
Rate Limit	Specify the rate limit value for Dynamic ARP Inspection. If the incoming rate exceeds the Rate Limit value for consecutively burst interval seconds, ARP packets will be dropped. If the value is **None**, there is no limit. The default is 15 packets per second (pps).
Burst Interval	Specify the burst interval for rate limiting on this interface. If the Rate Limit is **None,** the Burst Interval has no meaning and shows as **N/A** (Not Applicable). The default is 1 second.

- Click **Submit** to apply the new configuration and cause the change to take effect. These changes will not be retained across a power cycle unless a Save configuration is performed.
- Click **Refresh** to refresh the page with the most current data from the switch.

3.11.4 DAI ARP ACL Configuration

Use the DAI ARP ACL Configuration page to add or remove DAI ARP ACLs.

To display the DAI ARP ACL Configuration page, click **Switching > Dynamic ARP Inspection > DAI ARP ACL Configuration** in the navigation tree.

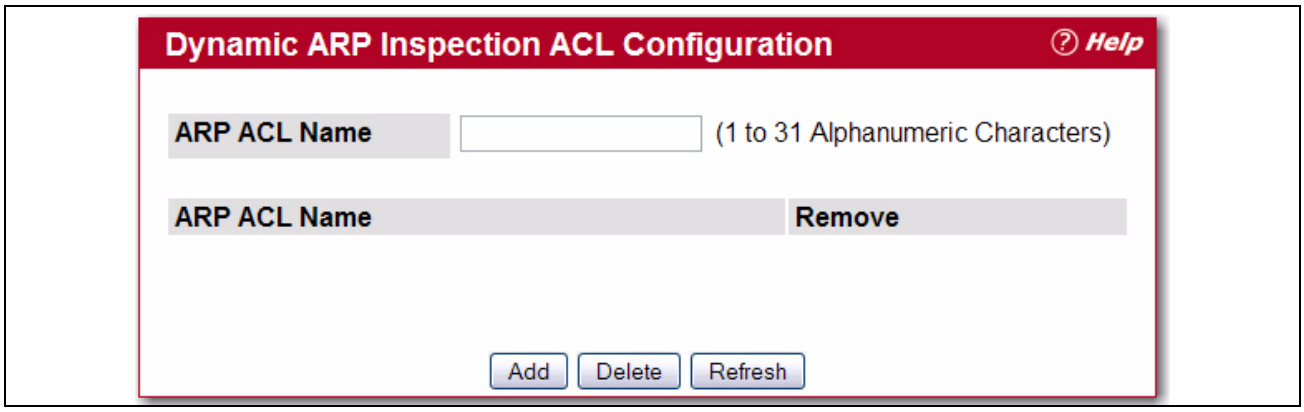

Figure 3-39: Dynamic ARP Inspection ARP ACL Configuration

Table 3-37: Dynamic ARP Inspection ARP ACL Configuration

Field	Description
ARP ACL Name	Use this field to create a new ARP ACL for Dynamic ARP Inspection. The name can be 1 to 31 alphanumeric characters in length.
ARP ACL List	Displays by name a list of all the configured ARP ACLs. Use the **Remove** column, to select the particular ACLs you want to delete.

- Click **Add** to create a new ARP ACL.

- Click **Delete** to remove the configured ARP ACL entry you selected in the **Remove** column.
- Click **Refresh** to refresh the page with the most current data from the switch.

3.11.5 DAI ARP ACL Rule Configuration

Use the DAI ARP ACL Rule Configuration page to add or remove DAI ARP ACL Rules.

To display the DAI ARP ACL Rule Configuration page, click **Switching > Dynamic ARP Inspection > DAI ARP ACL Rule Configuration** in the navigation tree.

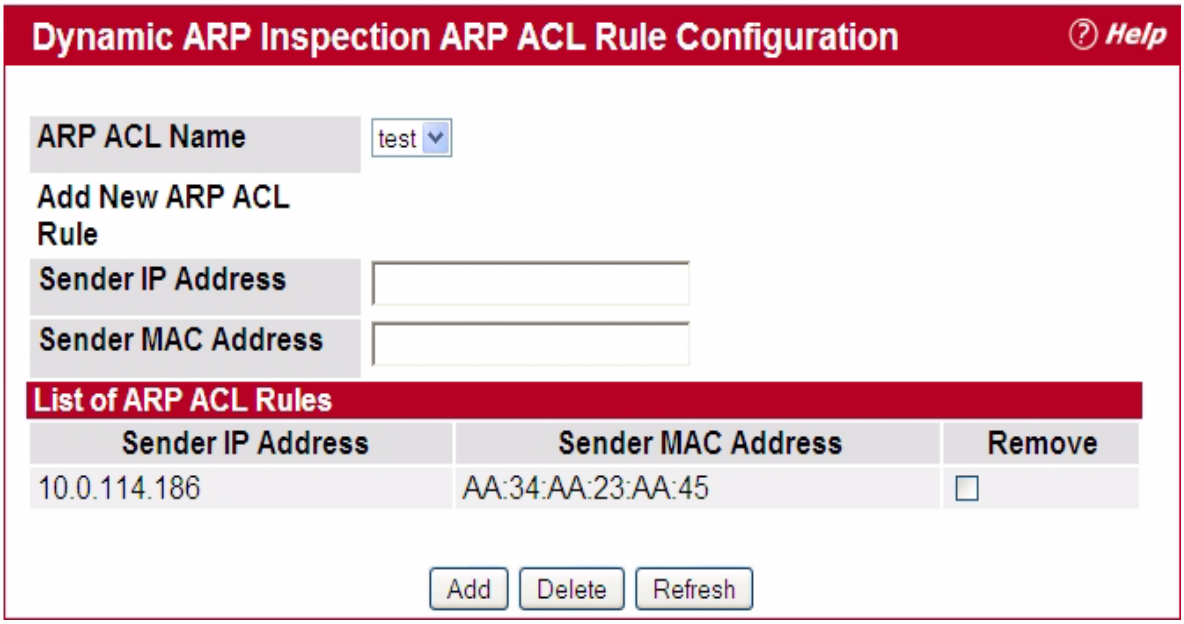

Figure 3-40: Dynamic ARP Inspection ARP ACL Rule Configuration

Table 3-38: Dynamic ARP Inspection ARP ACL Rule Configuration

Field	Description
ARP ACL Name	Select the ARP ACL for which information is to be displayed or configured.
Sender IP Address	To create a new rule for the selected ARP ACL, enter in this field the Sender IP Address match value for the ARP ACL.
Sender MAC Address	To create a new rule for the selected ARP ACL, enter in this field the Sender MAC Address match value for the ARP ACL.
Remove	Use the **Remove** column to select the particular ARP ACL Rules you want to delete.

- Click **Add** to add a new ARP ACL rule.
- Click **Submit** to delete the entries selected in the **Remove** column.
- Click **Refresh** to refresh the page with the most current data from the switch.

3.11.6 DAI Statistics

Use the DAI Statistics page to display the statistics per VLAN.

To display the DAI Statistics page, click **Switching > Dynamic ARP Inspection > DAI Statistics** in the navigation tree.

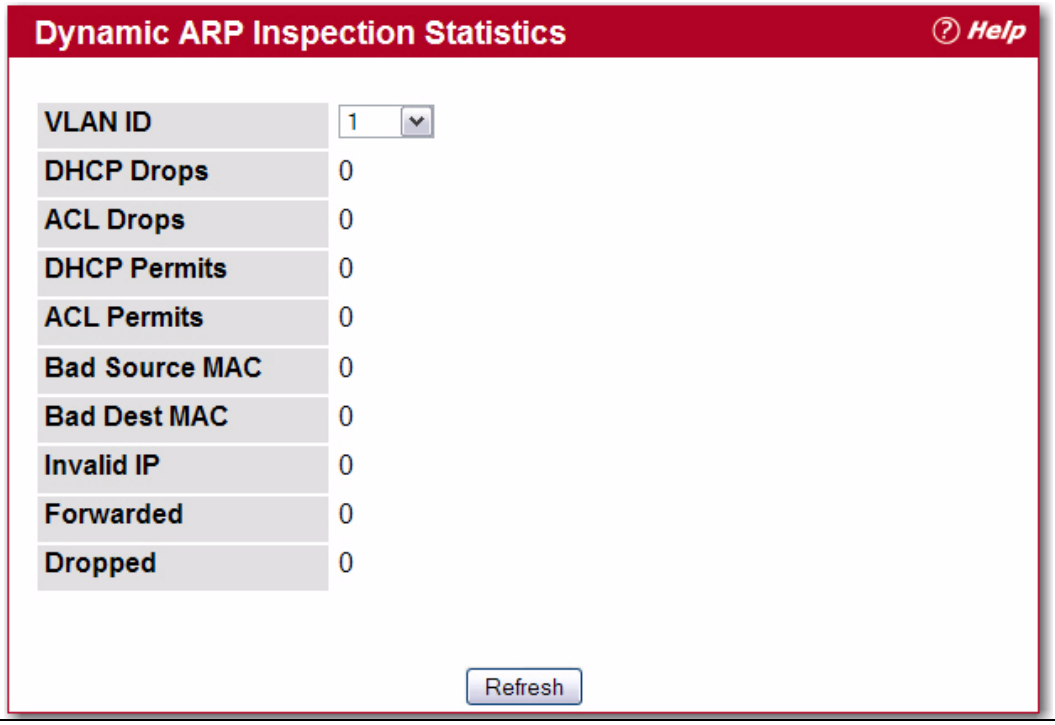

Figure 3-41: Dynamic ARP Inspection Statistics

Table 3-39: Dynamic ARP Inspection Statistics

Field	Description
VLAN ID	Select the DAI-enabled VLAN ID for which to display statistics.
DHCP Drops	The number of ARP packets that were dropped by DAI because there was no matching DHCP snooping binding entry found.
ACL Drops	The number of ARP packets that were dropped by DAI because there was no matching ARP ACL rule found for this VLAN and the static flag is set on this VLAN.
DHCP Permits	The number of ARP packets that were forwarded by DAI because there was a matching DHCP snooping binding entry found.
ACL Permits	The number or ARP packets that were permitted by DAI because there was a matching ARP ACL rule found for this VLAN.
Bad Source MAC	The number of ARP packets that were dropped by DAI because the sender MAC address in the ARP packet did not match the source MAC in the Ethernet header.
Bad Dest MAC	The number of ARP packets that were dropped by DAI because the target MAC address in the ARP reply packet did not match the destination MAC in the Ethernet header.
Invalid IP	The number of ARP packets that were dropped by DAI because the sender IP address in the ARP packet or target IP address in the ARP reply packet is not valid. Not valid addresses include 0.0.0.0, 255.255.255.255, IP multicast addresses, class E addresses (240.0.0.0/4), and loopback addresses (127.0.0.0/8).
Forwarded	The number of valid ARP packets forwarded by DAI.
Dropped	The number of not valid ARP packets dropped by DAI.

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

3.12 Configuring IGMP Snooping

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

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

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

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

3.12.1 Global Configuration and Status

Use the IGMP Snooping Global Configuration and Status page to enable IGMP snooping on the switch and view information about the current IGMP configuration.

To access the IGMP Snooping Configuration and Status page, click **Switching > IGMP Snooping > Configuration and Status** in the navigation tree.

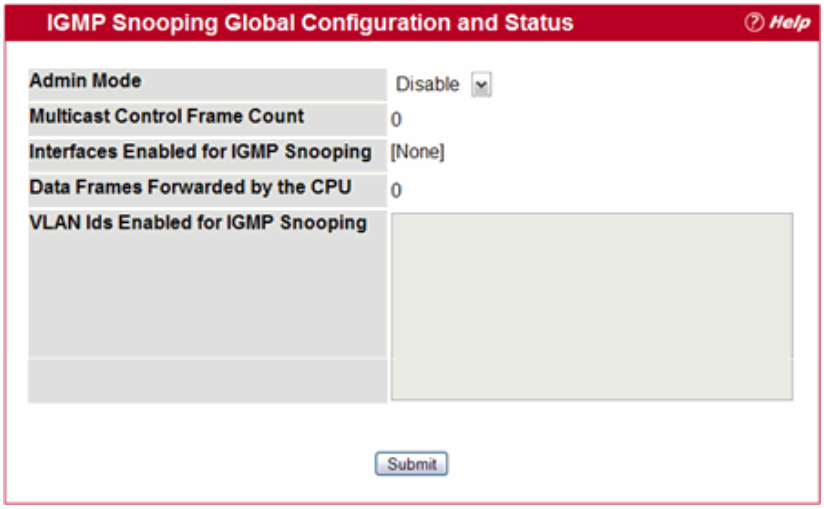

Figure 3-42: IGMP Snooping Global Configuration and Status

Table 3-40: IGMP Snooping Global Configuration and Status Fields

Field	Description
Admin Mode	Select the administrative mode for IGMP Snooping for the switch from the pulldown menu. The default is disable.
Multicast Control Frame Count	Shows the number of multicast control frames that have been processed by the CPU.
Interfaces Enabled for IGMP Snooping	Lists the interfaces currently enabled for IGMP Snooping. To enable interfaces for IGMP snooping, see 3.12.2 Interface Configuration.
Data Frames Forwarded by the CPU	Shows the number of data frames forwarded by the CPU.
VLAN Ids Enabled For IGMP Snooping	Displays VLAN Ids enabled for IGMP snooping. To enable VLANs for IGMP snooping, see 3.12.4 VLAN Configuration.

Select **Enable** or **Disable** the **Admin Mode** field and click **Submit** to turn the feature on or off. Perform a save if you want the changes to remain in effect over a power cycle.

3.12.2 Interface Configuration

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

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

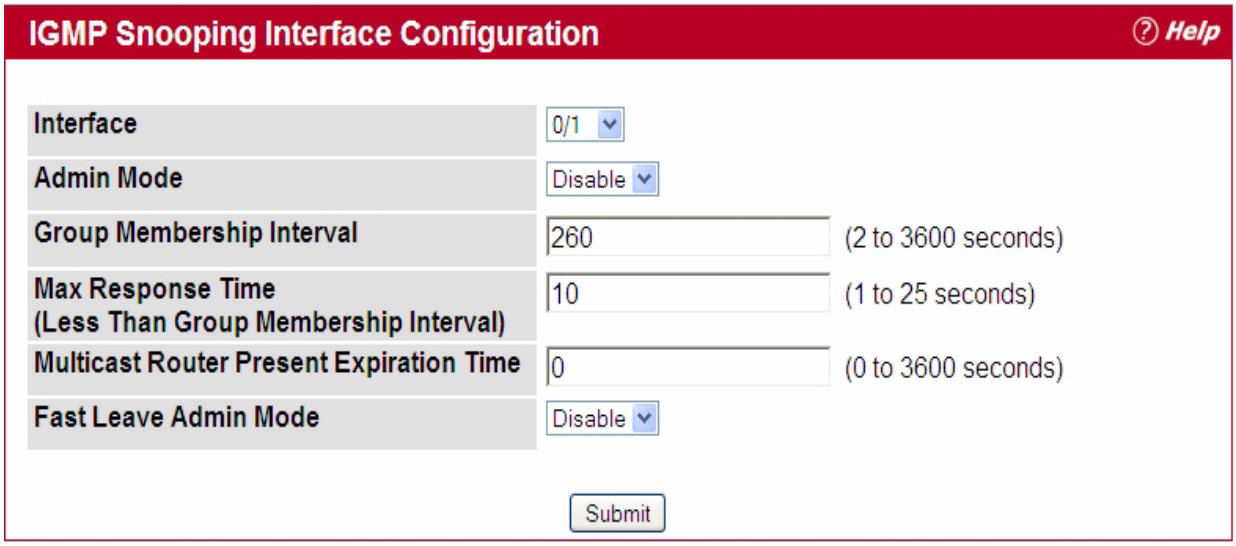

Figure 3-43: IGMP Snooping Interface Configuration

Table 3-41: IGMP Snooping Interface Configuration Fields

Field	Description
Slot/Port	Select the physical or LAG interfaces to configure.
Admin Mode	Select the interface mode for the selected interface for IGMP Snooping for the switch from the pulldown menu. The default is disable.
Group Membership Interval	Specify the amount of time you want the switch to wait for a report for a particular group on a particular interface before it deletes that interface from the group. The valid range is from (2 to 3600) seconds. The default is 260 seconds.
Max Response Time	Specify the amount of time you want the switch to wait after sending a query on an interface because it did not receive a report for a particular group on that interface. Enter a value greater or equal to 1 and less than the Group Membership Interval in seconds. The default is 10 seconds. The configured value must be less than the Group Membership Interval.
Multicast Router Present Expiration Time	Specify the amount of time you want the switch to wait to receive a query on an interface before removing it from the list of interfaces with multicast routers attached. Enter a value between 0 and 3600 seconds. The default is 0 seconds. A value of zero indicates an infinite timeout; i.e., no expiration.
Fast Leave Admin Mode	Select the Fast Leave mode for the a particular interface from the pulldown menu. The default is **Disable**.

If you make any changes on the page, click **Submit** to apply the new settings to the switch.

3.12.3 VLAN Status

Use the IGMP Snooping VLAN Status page to view information about the VLANs on the system that are configured for IGMP snooping.

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

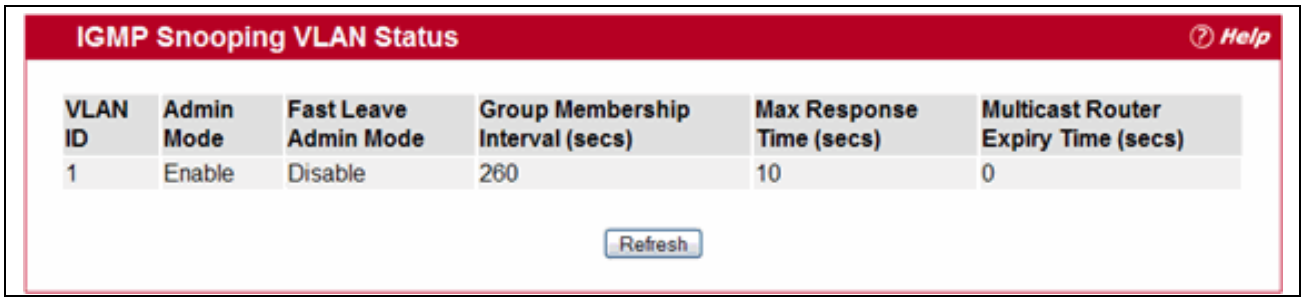

VLAN ID	Admin Mode	Fast Leave Admin Mode	Group Membership Interval (secs)	Max Response Time (secs)	Multicast Router Expiry Time (secs)
1	Enable	Disable	260	10	0

Figure 3-44: IGMP Snooping VLAN Status

Table 3-42: IGMP Snooping VLAN Status Fields

Field	Description
VLAN ID	Displays the VLAN IDs for which the IGMP Snooping mode is Enabled.
Admin Mode	Shows the IGMP Snooping Mode for the VLAN ID.
Fast Leave Admin Mode	Indicates whether IGMP Snooping Fast-leave is active on the VLAN.
Group Membership Interval	Shows the amount of time in seconds that a switch will wait for a report from a particular group on a particular interface, which is participating in the VLAN, before deleting the interface from the entry.
Maximum Response Time	Shows the amount of time the switch waits after it sends a query on an interface, participating in the VLAN, because it did not receive a report for a particular group on that interface. This value may be configured.
Operational Maximum Response Time	Displays the value for maximum response time of IGMP Snooping for the specified VLAN ID. Its value is learned dynamically from the IGMPv2 or IGMPv3 queries received on this VLAN.
Multicast Router Expiry Time	Shows the amount of time to wait before removing an interface that is participating in the VLAN from the list of interfaces with multicast routers attached. The interface is removed if a query is not received.

Click **Refresh** to re-display the page with the latest information from the router.

3.12.4 VLAN Configuration

Use the IGMP Snooping VLAN Configuration page to configure IGMP snooping settings for VLANs on the system.

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

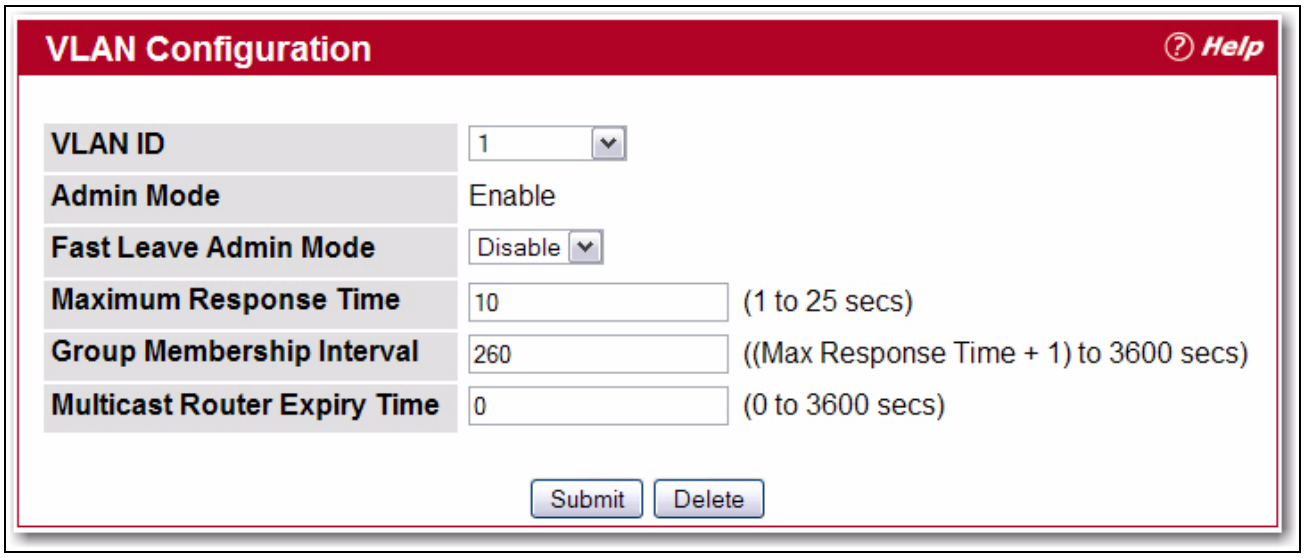

Figure 3-45: IGMP Snooping VLAN Configuration

Table 3-43: IGMP Snooping VLAN Configuration Fields

Field	Description
VLAN ID	From the drop-down menu, select the VLAN ID of the VLAN to modify, or select New Entry to configure settings for a VLAN that does not have IGMP Snooping enabled.
Admin Mode	Enable is the only available option from the drop-down menu. To disable the IGMP snooping admin mode on the VLAN, select the VLAN from the VLAN ID field and click **Delete**.
Fast Leave Admin Mode	Enabling fast-leave allows the switch to immediately remove the layer 2 LAN interface from its forwarding table entry upon receiving an IGMP leave message for that multicast group without first sending out MAC-based general queries to the interface. You should enable fast-leave admin mode only on VLANs where only one host is connected to each layer 2 LAN port. This prevents the inadvertent dropping of the other hosts that were connected to the same layer 2 LAN port but were still interested in receiving multicast traffic directed to that group. Also, fast-leave processing is supported only with IGMP version 2 hosts.
Group Membership Interval	The Group Membership Interval time is the amount of time in seconds that a switch waits for a report from a particular group on a particular interface before deleting the interface from the entry. This value must be greater than the IGMPv3 Maximum Response time value. The range is 2 to 3600 seconds.
Maximum Response Time	Enter the amount of time in seconds that a switch will wait after sending a query on an interface because it did not receive a report for a particular group in that interface. This value must be less than the IGMP Query Interval time value. The range is 1 to 3599 seconds.
Operational Maximum Response Time	This read-only field displays the value for maximum response time of IGMP Snooping for the specified VLAN ID. Its value is learned dynamically from the IGMPv2 or IGMPv3 queries received on this VLAN. For the multicast traffic not to get disturbed, you should configure group membership interval to be greater than this value.
Multicast Router Expiry Time	Enter the amount of time in seconds that a switch waits for a query to be received on an interface before the interface is removed from the list of interfaces with multicast routers attached. The range is 0 to 3600 seconds. A value of 0 indicates an infinite time-out; i.e., no expiration.

If you make any changes to the page, click **Submit** to apply the new settings to the system.

3.12.5 Multicast Router Status

Use the IGMP Snooping Multicast Router Status page to see whether a particular interface is configured as a multicast router interface.

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

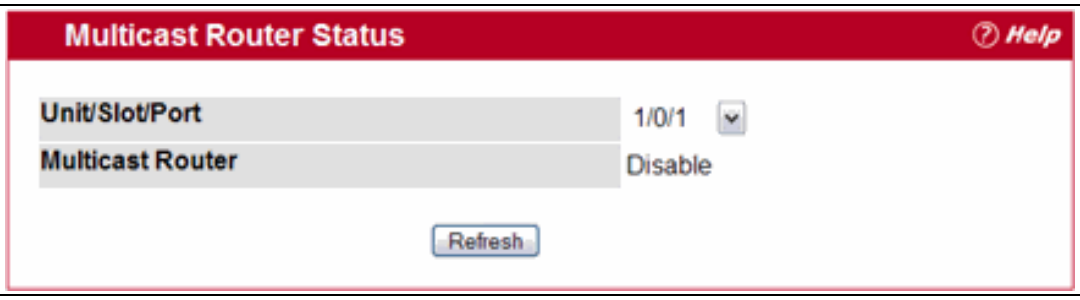

Figure 3-46: Multicast Router Status

Table 3-44: Multicast Router Status Fields

Field	Description
Slot/Port	Select the physical or LAG interface to display.
Multicast Router	Shows whether the specified interface is configured as a multicast router interface.

Click **Refresh** to re-display the page with the latest information from the router.

3.12.6 Multicast Router Configuration

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

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

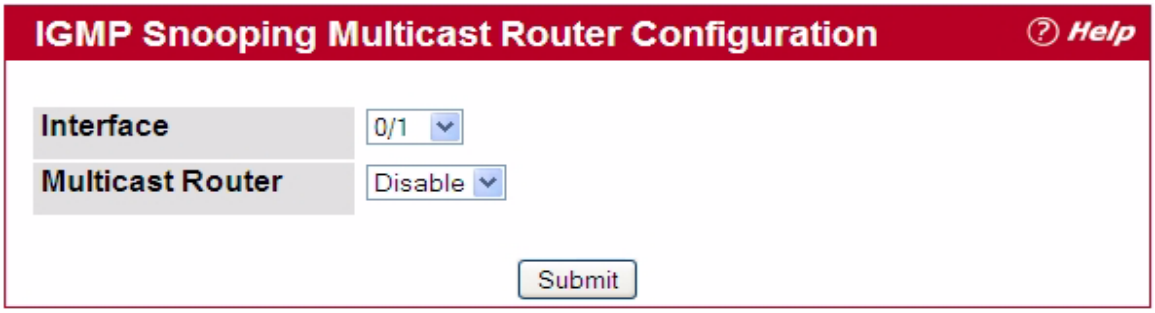

Figure 3-47: Multicast Router Configuration

Table 3-45: Multicast Router Configuration Fields

Field	Description
Slot/Port	Select the physical or LAG interface to display.
Multicast Router	Set the multicast router status: • **Enabled**: The port is a multicast router interface. • **Disabled**: The port does not have a multicast router configured.

If you enable or disable multicast router configuration on an interface, click **Submit** to apply the new settings to the switch.

3.12.7 Multicast Router VLAN Status

Use the IGMP Snooping Multicast Router VLAN Status page to view multicast router settings for VLANs on a specific interface.

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

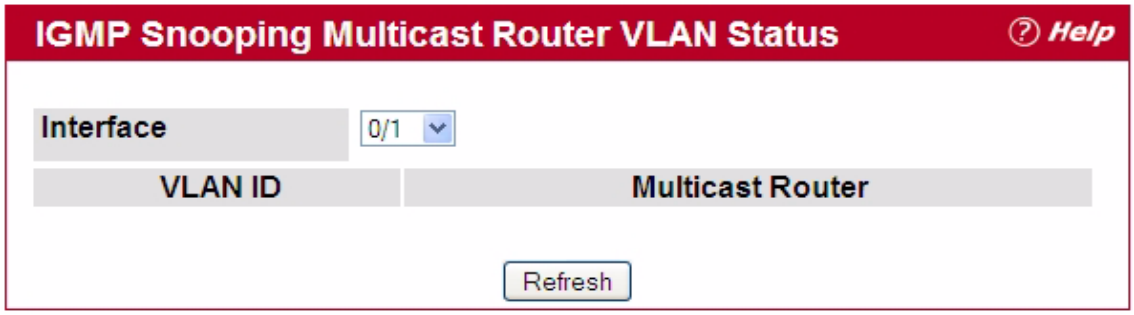

Figure 3-48: Multicast Router VLAN Status

The IGMP Snooping Multicast Router VLAN Status page contains the following fields:

Table 3-46: Multicast Router VLAN Status Fields

	Description
Slot/Port	Select the physical or LAG interface to display.
VLAN ID	If a VLAN is enabled for multicast routing on the interface, this field displays its ID.
Multicast Router	Indicates that the multicast router is enabled for the VLAN on this interface.

Click **Refresh** to re-display the page with the latest information from the router.

3.12.8 Multicast Router VLAN Configuration

Use the IGMP Snooping Multicast Router VLAN Configuration page to configure multicast router settings for VLANs on an interface.

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

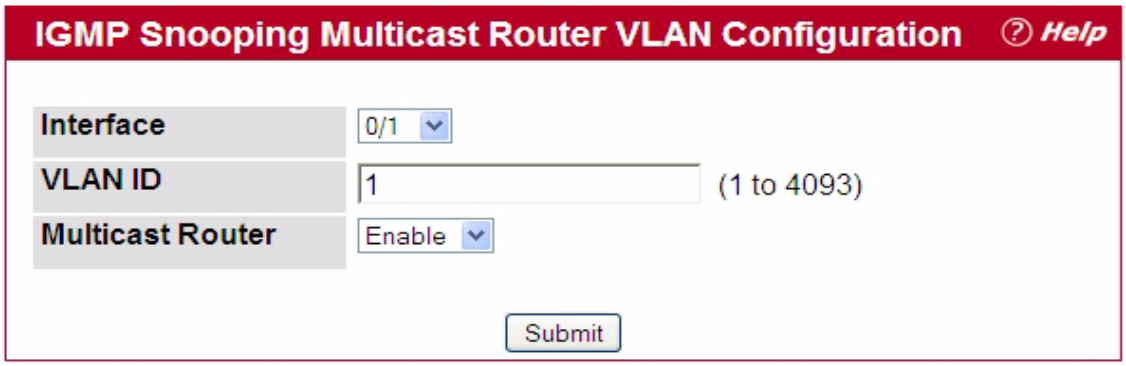

Figure 3-49: Multicast Router VLAN Configuration

Table 3-47: Multicast Router VLAN Configuration Fields

Field	Description
Slot/Port	Select the physical or LAG interface to display.
VLAN ID	Enter the VLAN ID to configure as enabled or disabled for multicast routing.
Multicast Router	Select Enable or Disable from the drop-down menu to change the multicast router mode of the VLAN associated with this interface.

If you enable or disable multicast router configuration for VLANs on an interface, click **Submit** to apply the new settings to the switch.

3.13 Configuring IGMP Snooping Queriers

IGMP snooping requires that one central switch or router periodically query all end-devices on the network to announce their multicast memberships. This central device is the 'IGMP querier'. The IGMP query responses, known as IGMP reports, keep the switch updated with the current multicast group membership on a port-by-port basis. If the switch does not receive updated membership information in a timely fashion, it will stop forwarding multicasts to the port where the end device is located.

These pages enable you to configure and display information on IGMP snooping queriers on the network and, separately, on VLANs.

3.13.1 IGMP Snooping Querier Configuration

Use this page to enable or disable the IGMP Snooping Querier feature, specify the IP address of the router to perform the querying, and configure related parameters. Users must have Read/Write access privileges to change the data on this page.

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

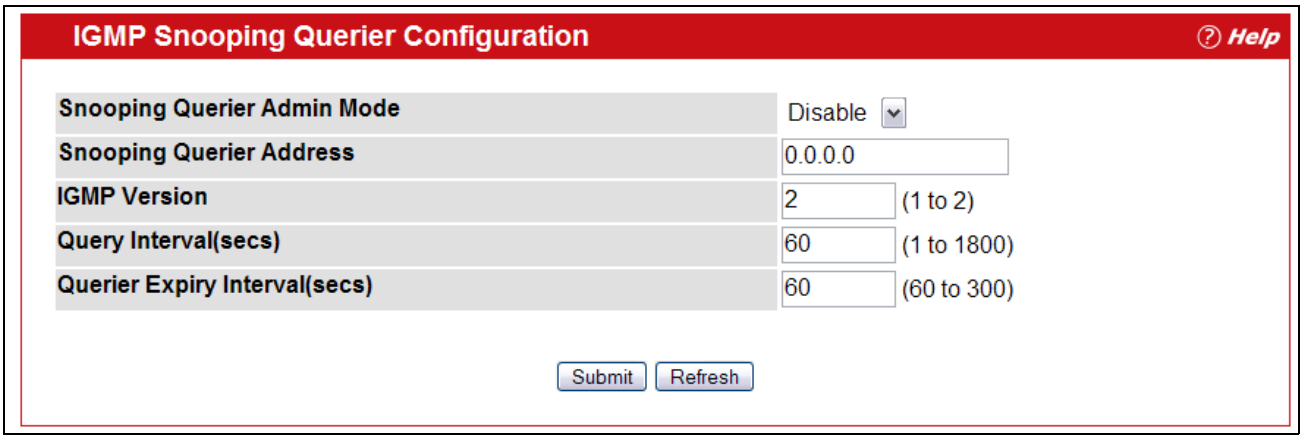

Figure 3-50: IGMP Snooping Querier Configuration

Table 3-48: IGMP Snooping Querier Configuration Fields

Field	Description
Snooping Querier Admin Mode	Select the administrative mode for IGMP Snooping for the switch from the pulldown menu. The default is **Disable**.
Snooping Querier Address	Specify the Snooping Querier Address to be used as source address in periodic IGMP queries. This address is used when no address is configured on the VLAN on which query is being sent.
IGMP Version	Specify the IGMP protocol version used in periodic IGMP queries.
Query Interval	Specify the time interval in seconds between periodic queries sent by the snooping querier. The Query Interval must be a value in the range of 1 and 1800. The default value is 60.
Querier Expiry Interval	Specify the time interval in seconds after which the last querier information is removed. The Querier Expiry Interval must be a value in the range of 60 and 300. The default value is 60.

- If you configure an IGMP snooping querier, click **Submit** to apply the new settings to the switch.
- Click **Refresh** to re-display the page with the latest information from the switch.

3.13.2 IGMP Snooping Querier VLAN Configuration

Use this page to configure IGMP queriers for use with VLANs on the network.

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

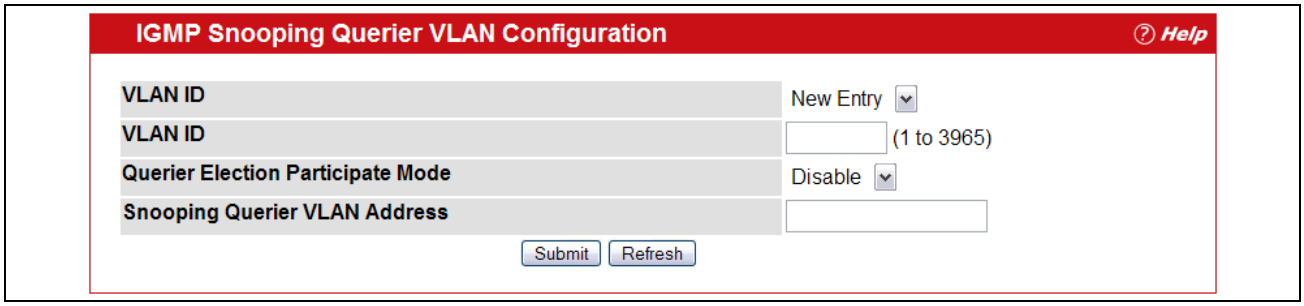

Figure 3-51: IGMP Snooping Querier VLAN Configuration

Table 3-49: IGMP Snooping Querier VLAN Configuration Fields

Field	Description
VLAN ID	Specifies VLAN ID for which the IGMP Snooping Querier is to be enabled. Select **New Entry** to create a new VLAN ID for IGMP Snooping.
Querier Election Participate Mode	Enables or disables Querier Participate Mode. When this mode is disabled, upon seeing another querier of same version in the VLAN, the snooping querier moves to non-querier state. When enabled, the snooping querier participates in querier election, in which the least IP address operates as the querier in that VLAN. The other querier moves to non-querier state.
Snooping Querier VLAN Address	Specifies the Snooping Querier Address to be used as source address in periodic IGMP queries sent on the specified VLAN.

- If you configure a snooping querier for a VLAN, click **Submit** to apply the new settings.
- Click **Refresh** to re-display the page with the latest information from the switch.

3.13.3 IGMP Snooping Querier VLAN Configuration Summary

Use this page to view summary information for IGMP snooping queriers for on VLANs in the network.

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

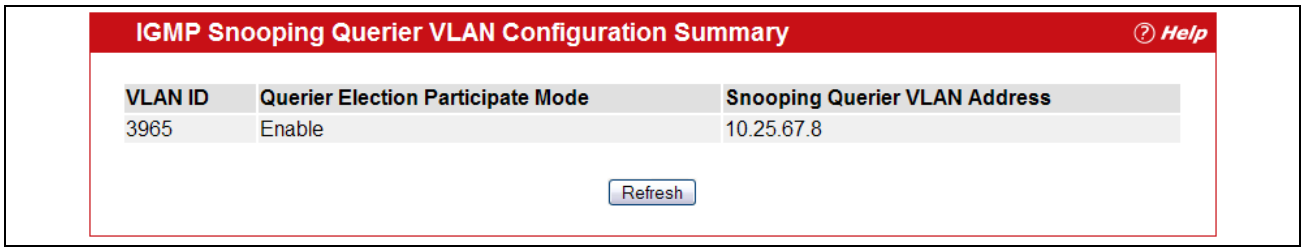

Figure 3-52: IGMP Snooping Querier VLAN Configuration Summary

Table 3-50: IGMP Snooping Querier VLAN Configuration Summary Fields

Field	Description
VLAN ID	Specifies the VLAN ID on which IGMP Snooping Querier is administratively enabled.
Querier Election Participate Mode	Displays the querier election participate mode on the VLAN. When this mode is disabled, up on seeing a query of the same version in the VLAN, the snooping querier moves to non-querier state. When this mode is enabled, the snooping querier participate in querier election, in which the lowest IP address operates as the querier in that VLAN. The other querier moves to non-querier state.
Snooping Querier VLAN Address	Displays the Snooping Querier Address to be used as source address in periodic IGMP queries sent on the specified VLAN.

Click **Refresh** to re-display the page with the latest information from the router.

3.13.4 IGMP Snooping Querier VLAN Status

Use this page to view the operational state and other information for IGMP snooping queriers for VLANs on the network.

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

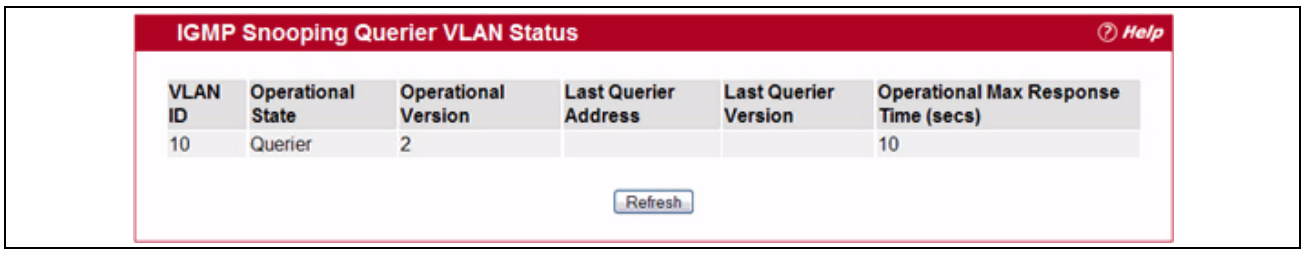

Figure 3-53: IGMP Snooping Querier VLAN Status

Table 3-51: IGMP Snooping Querier VLAN Status Fields

Field	Description
VLANID	Specifies the VLAN ID on which the IGMP Snooping Querier is administratively enabled and for which VLAN exists in the VLAN database.
Operational State	Specifies the operational state of the IGMP Snooping Querier on a VLAN: • **Querier**: The snooping switch is the querier in the VLAN. The snooping switch will send out periodic queries with a time interval equal to the configured querier query interval. If the snooping switch sees a better querier (numerically lower) in the VLAN, it moves to non-querier mode. • **Non-Querier**: The snooping switch is in non-querier mode in the VLAN. If the querier expiry interval timer expires, the snooping switch moves into querier mode. • **Disabled**: The snooping querier is not operational on the VLAN. The snooping querier moves to disabled mode when IGMP snooping is not operational on the VLAN, when the querier address is not configured, or the network management address is not configured.
Operational Version	Displays the IGMP protocol version of the operational querier.
Last Querier Address	Displays the IP address of the last querier from which a query was snooped on the VLAN.
Last Querier Version	Displays the IGMP protocol version of the last querier from which a query was snooped on the VLAN.
Operational Max Response Time	Displays the maximum response time to be used in the queries that are sent by the snooping querier.

Click **Refresh** to re-display the page with the latest information from the switch.

3.14 Configuring MLD Snooping

In IPv4, Layer 2 switches can use IGMP snooping to limit the flooding of multicast traffic by dynamically configuring Layer-2 interfaces so that multicast traffic is forwarded to only those interfaces associated with an IP multicast address. In IPv6, MLD snooping performs a similar function. With MLD snooping, IPv6 multicast data is selectively forwarded to a list of ports that want to receive the data, instead of being flooded to all ports in a VLAN. This list is constructed by snooping IPv6 multicast control packets.

MLD is a protocol used by IPv6 multicast routers to discover the presence of multicast listeners (nodes wishing to receive IPv6 multicast packets) on its directly-attached links and to discover which multicast packets are of interest to neighboring nodes. MLD is derived from IGMP; MLD version 1 (MLDv1) is equivalent to IGMPv2 and MLD version 2 (MLDv2) is equivalent to IGMPv3. MLD is a subprotocol of Internet Control Message Protocol version 6 (ICMPv6), and MLD messages are a subset of ICMPv6 messages.

The switch can snoop on both MLDv1 and MLDv2 protocol packets and bridge IPv6 multicast data based on destination IPv6 multicast MAC addresses. The switch can be configured to perform MLD snooping and IGMP snooping simultaneously.

3.14.1 Configuration and Status

Use the MLD Snooping Global Configuration and Status page to enable MLD snooping on the switch and view information about the current MLD snooping configuration.

To access this page, click **Switching > MLD Snooping > Configuration and Status** in the navigation tree.

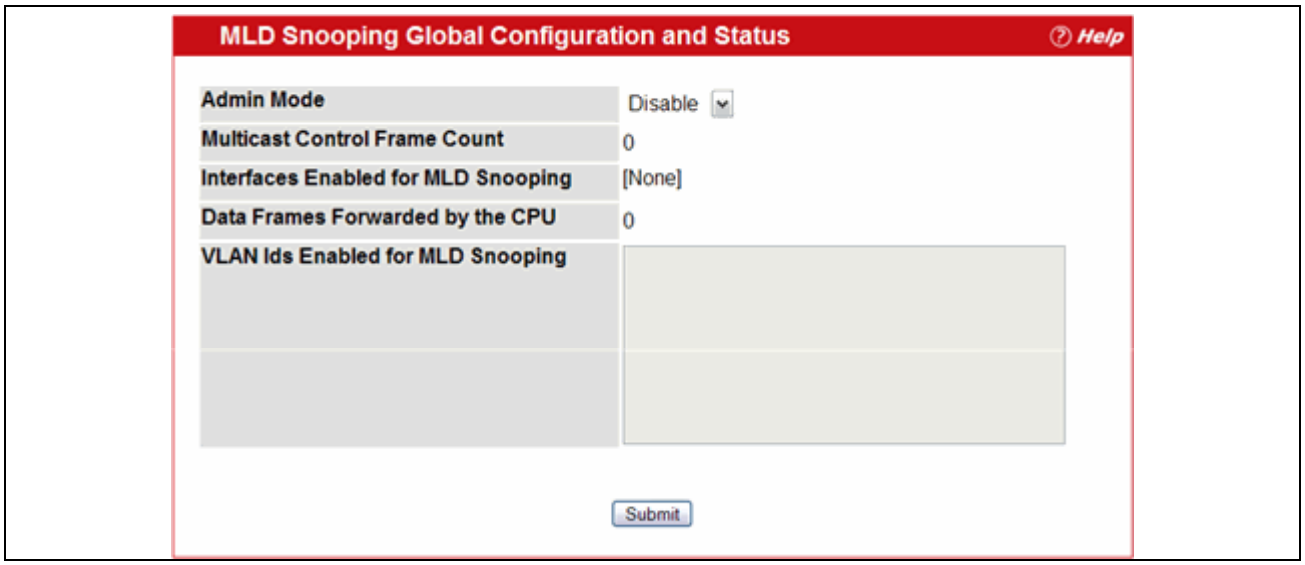

Figure 3-54: MLD Snooping Global Configuration and Status

Table 3-52: MLD Snooping Global Configuration and Status Fields

Field	Description
Admin Mode	Select the administrative mode for MLD Snooping for the switch from the pulldown menu. The default is disable.
Multicast Control Frame Count	Shows the number of multicast control frames that have been processed by the CPU.
Interfaces Enabled for MLD Snooping	Lists the interfaces currently enabled for MLD Snooping. To enable interfaces for MLD snooping, see 3.14.2 Interface Configuration.
Data Frames Forwarded by the CPU	Shows the number of data frames forwarded by the CPU.
VLAN Ids Enabled For MLD Snooping	Displays VLAN Ids enabled for MLD snooping. To enable interfaces for MLD snooping, see 3.14.4 VLAN Configuration.

Select **Enable** or **Disable** the **Admin Mode** field and click **Submit** to turn the feature on or off. Perform a save if you want the changes to remain in effect over a power cycle.

3.14.2 Interface Configuration

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

To access the MLD Snooping Interface Configuration page, click **Switching > MLD Snooping > Interface Configuration** in the navigation tree.

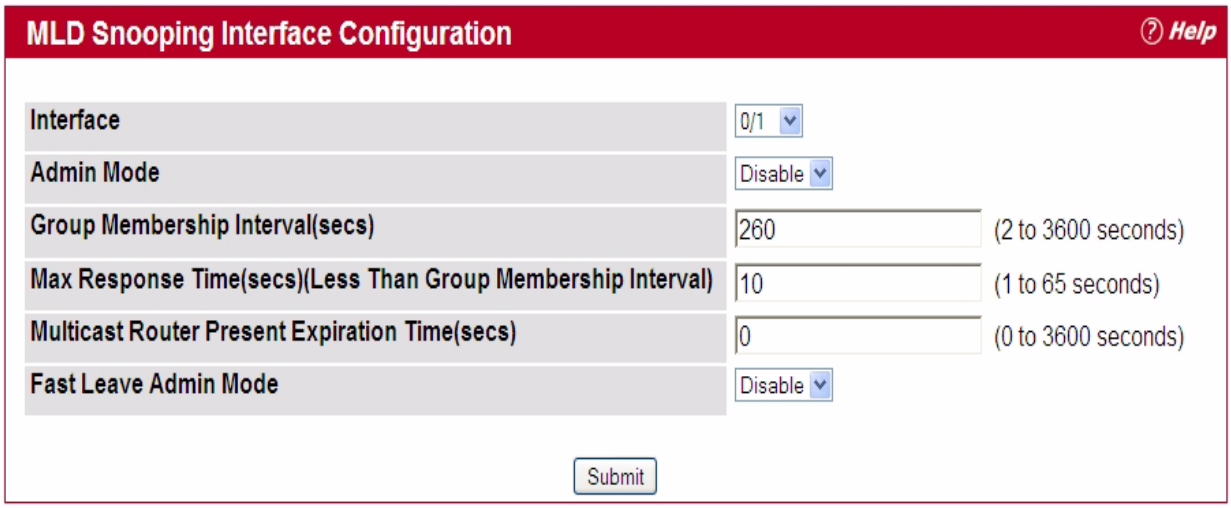

Figure 3-55: MLD Snooping Interface Configuration

Table 3-53: MLD Snooping Interface Configuration Fields

Field	Description
Slot/Port	Select the physical or LAG interfaces to configure.
Admin Mode	Select the interface mode for the selected interface for MLD Snooping for the switch from the pulldown menu. The default is **Disable**.
Group Membership Interval	Specify the amount of time you want the switch to wait for a report for a particular group on a particular interface before it deletes that interface from the group. The valid range is from (2 to 3600) seconds. The default is 260 seconds.
Max Response Time	Specify the amount of time you want the switch to wait after sending a query on an interface because it did not receive a report for a particular group on that interface. Enter a value greater or equal to 1 and less than the Group Membership Interval in seconds. The default is 10 seconds. The configured value must be less than the Group Membership Interval.
Multicast Router Present Expiration Time	Specify the amount of time you want the switch to wait to receive a query on an interface before removing it from the list of interfaces with multicast routers attached. Enter a value between 0 and 3600 seconds. The default is 0 seconds. A value of zero indicates an infinite timeout; i.e., no expiration.
Fast Leave Admin Mode	Select the Fast Leave mode for the a particular interface from the pulldown menu. The default is **Disable**.

If you make any changes on the page, click **Submit** to apply the new settings to the switch.

3.14.3 VLAN Status

Use the MLD Snooping VLAN Status page to view information about the VLANs on the system that are configured for MLD snooping.

To access the MLD Snooping VLAN Status page, click **Switching > MLD Snooping > VLAN Status** in the navigation tree.

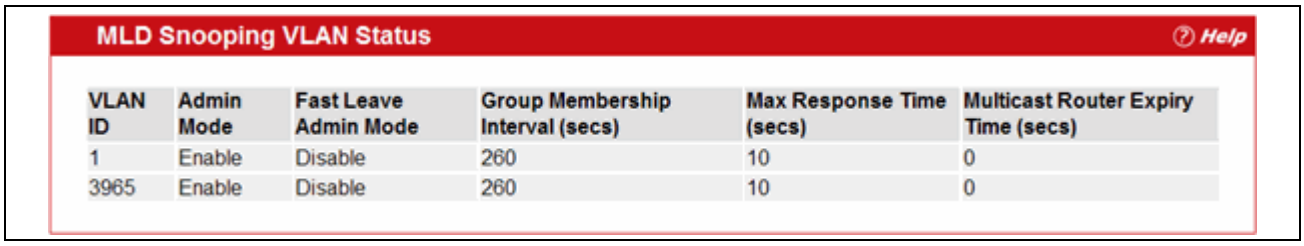

VLAN ID	Admin Mode	Fast Leave Admin Mode	Group Membership Interval (secs)	Max Response Time (secs)	Multicast Router Expiry Time (secs)
1	Enable	Disable	260	10	0
3965	Enable	Disable	260	10	0

Figure 3-56: MLD Snooping VLAN Status

Table 3-54: MLD Snooping VLAN Status Fields

Field	Description
VLAN ID	Displays the VLAN IDs for which the MLD Snooping mode is Enabled.
Admin Mode	Shows the MLD Snooping Mode for the VLAN ID.
Fast Leave Admin Mode	Indicates whether MLD Snooping Fast-leave is active on the VLAN.
Group Membership Interval	Shows the amount of time in seconds that a switch will wait for a report from a particular group on a particular interface, which is participating in the VLAN, before deleting the interface from the entry. The valid range is 2 to 3600.
Maximum Response Time	Shows the amount of time the switch waits after it sends a query on an interface, participating in the VLAN, because it did not receive a report for a particular group on that interface. The valid range is 1 to 3599. Its value should be greater than group membership interval value.
Multicast Router Expiry Time	Shows the amount of time to wait before removing an interface that is participating in the VLAN from the list of interfaces with multicast routers attached. The interface is removed if a query is not received. The valid range is 0 to 3600.

Click **Refresh** to re-display the page with the latest information from the router.

3.14.4 VLAN Configuration

Use the MLD Snooping VLAN Configuration page to configure MLD Snooping settings for VLANs on the system.

To access the MLD Snooping VLAN Configuration page, click **Switching > MLD Snooping > VLAN Configuration** in the navigation tree.

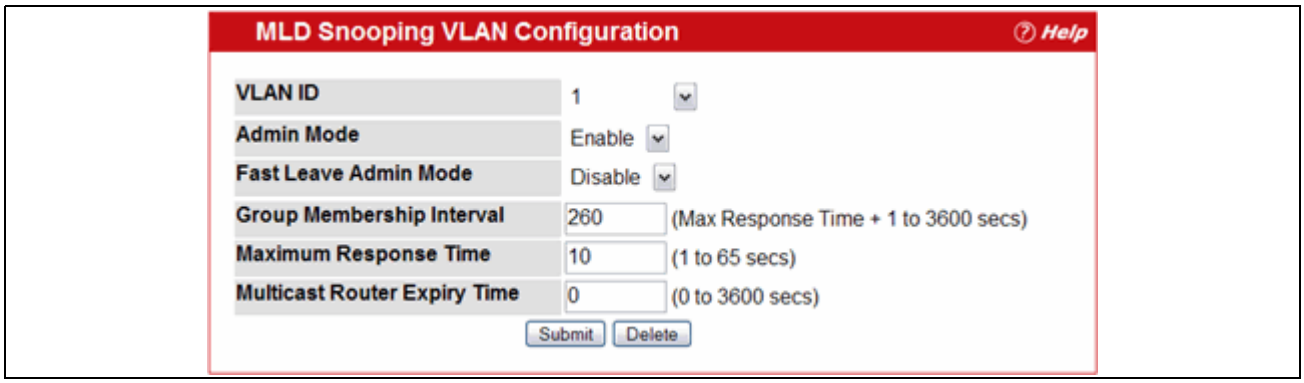

Figure 3-57: MLD Snooping VLAN Configuration

Table 3-55: MLD Snooping VLAN Configuration Fields

Field	Description
VLAN ID	Specifies list of VLAN IDs for which MLD Snooping is enabled. If no entries exist, New Entry displays. Enter the VLAN ID of the VLAN on which to enable and configure MLD Snooping.
Admin Mode	Enable is the only available option from the drop-down menu. To disable the MLD Snooping admin mode on the VLAN, select the VLAN from the VLAN ID field and click **Delete**.
Fast Leave Admin Mode	Enabling fast-leave allows the switch to immediately remove the layer-2 LAN interface from its forwarding table entry upon receiving an MLD leave message for that multicast group without first sending out MAC-based general queries to the interface. Enable fast-leave admin mode only on VLANs where only one host is connected to each layer-2 LAN port. This prevents the inadvertent dropping of the other hosts that were connected to the same layer-2 LAN port but were still interested in receiving multicast traffic directed to that group.
Group Membership Interval	The Group Membership Interval time is the amount of time in seconds that a switch waits for a report from a particular group on a particular interface before deleting the interface from the entry. This value must be greater than the Maximum Response time value. The range is 2 to 3600 seconds.
Maximum Response Time	Enter the amount of time in seconds that a switch will wait after sending a query on an interface because it did not receive a report for a particular group in that interface. This value must be less than the Group Membership Interval value. The range is 1 to 3599 seconds.
Multicast Router Expiry Time	Enter the amount of time in seconds that a switch waits for a query to be received on an interface before the interface is removed from the list of interfaces with multicast routers attached. The range is 0 to 3600 seconds. A value of 0 indicates an infinite time-out; i.e., no expiration.

- If you make any changes to the page, click **Submit** to apply the new settings to the system.
- To disable the MLD Snooping admin mode on a VLAN, select the VLAN from the VLAN ID field and click **Delete**.

3.14.5 Multicast Router Status

Use the MLD Snooping Multicast Router Status page to view multicast router functionality on selected ports. To access this page, click **Switching > MLD Snooping > Multicast Router Statistics** in the navigation tree.

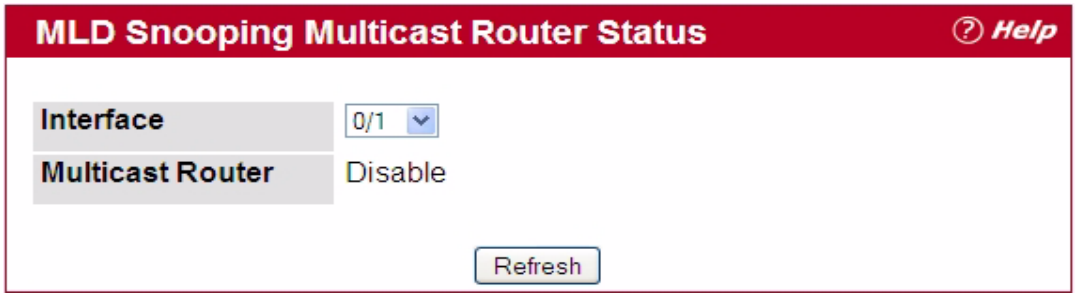

Figure 3-58: MLD Snooping Multicast Router Status

Table 3-56: MLD Snooping Multicast Router Status Fields

Field	Description
Slot/Port	Select the unit, slot and port number with the information to view.
Multicast Router	Indicates whether the specified interface is configured to perform multicast routing.

Click **Refresh** to re-display the page with the latest information from the router.

3.14.6 Multicast Router Configuration

The switch can dynamically learn of an attached multicast router, or you can configure a switch port as a multicast router interface. Use the MLD Snooping Multicast Router Configuration page to configure an interface as a static multicast router interface.

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

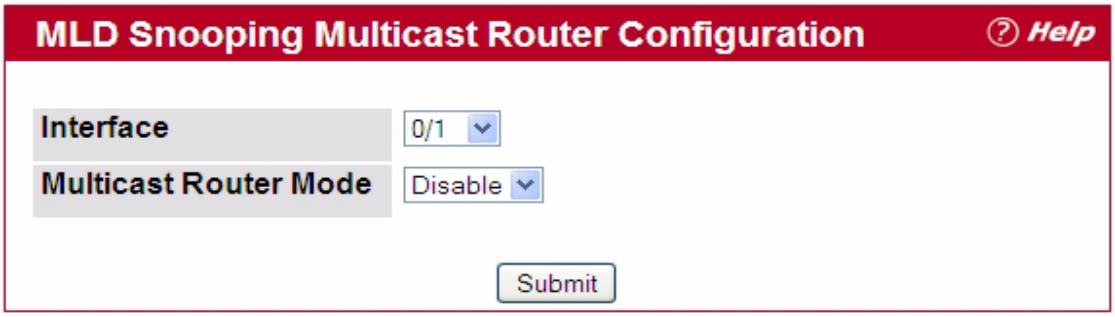

Figure 3-59: MLD Snooping Multicast Router Configuration

Table 3-57: MLD Snooping Multicast Router Configuration Fields

Field	Description
Slot/Port	Select the physical or LAG interface to display.
Multicast Router	Set the multicast router status: • **Enabled**: The port is a multicast router interface. • **Disabled**: The port does not have multicast router configured.

If you enable or disable multicast router configuration on an interface, click **Submit** to apply the new settings to the switch.

3.14.7 Multicast Router VLAN Status

Use the MLD Snooping Multicast Router VLAN Status page to view multicast router settings for VLANs on a specific interface.

To access the MLD Snooping Multicast Router VLAN Statistics page, click **Switching > MLD Snooping > Multicast Router VLAN Status** in the navigation tree.

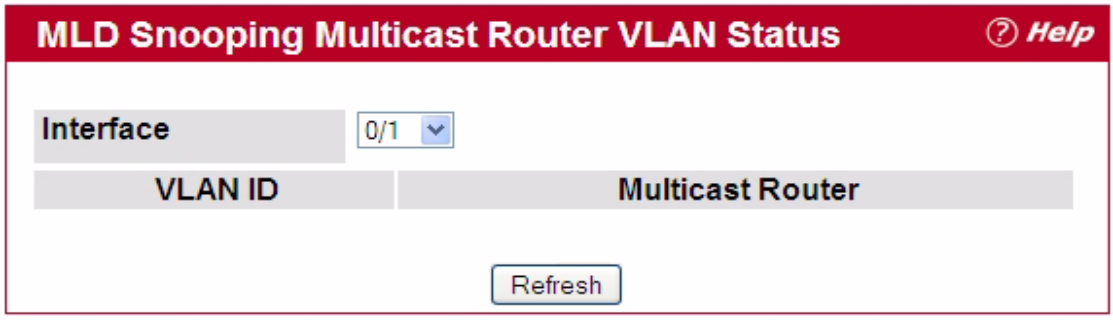

Figure 3-60: MLD Snooping Multicast Router VLAN Status

The MLD Snooping Multicast Router VLAN Statistics page contains the following fields:

Table 3-58: MLD Snooping Multicast Router VLAN Status Fields

	Description
Slot/Port	Select the physical or LAG interface to display.
VLAN ID	If a VLAN is enabled for multicast routing on the interface, this field displays its ID.
Multicast Router	Indicates that the multicast router is enabled for the VLAN on this interface.

Click **Refresh** to re-display the page with the latest information from the router.

3.14.8 Multicast Router VLAN Configuration

Use the MLD Snooping Multicast Router VLAN Configuration page to configure multicast router settings for VLANs on an interface.

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

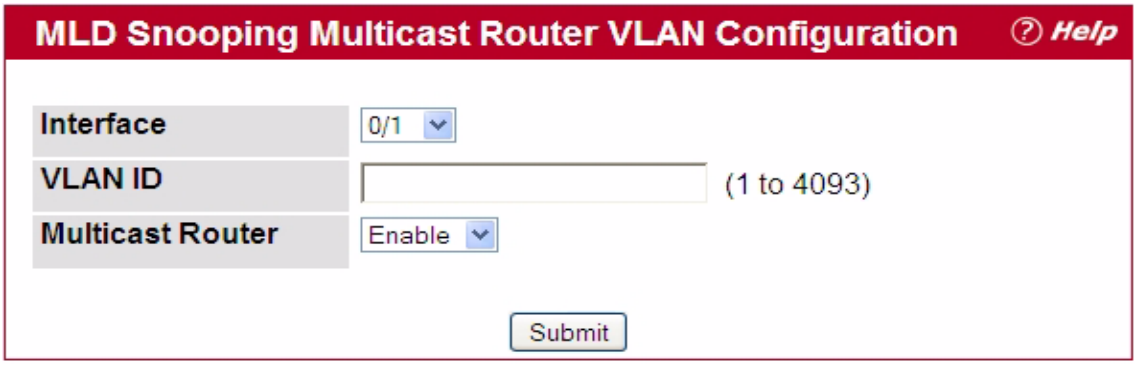

Figure 3-61: Multicast Router VLAN Configuration

Table 3-59: Multicast Router VLAN Configuration Fields

Field	Description
Slot/Port	Select the physical, VLAN, or LAG interface to display.
VLAN ID	Enter the VLAN ID to configure as enabled or disabled for multicast routing.
Multicast Router	Select Enable or Disable from the drop-down menu to change the multicast router mode of the VLAN associated with this interface.

If you enable or disable multicast router configuration for VLANs on an interface, click **Submit** to apply the new settings to the switch.

3.15 Configuring MLD Snooping Queriers

In an IPv6 environment, MLD Snooping requires that one central switch or router periodically query all end-devices on the network to announce their multicast memberships. This central device is the 'MLD querier'. The MLD query responses, known as MLD reports, keep the switch updated with the current multicast group membership on a port-by-port basis. If the switch does not receive updated membership information in a timely fashion, it will stop forwarding multicasts to the port where the end device is located.

These pages enable you to configure and display information on MLD Snooping queriers on the network and, separately, on VLANs.

3.15.1 MLD Snooping Querier Configuration

Use this page to enable or disable the MLD Snooping Querier feature, specify the IP address of the router to perform the querying, and configure related parameters. Users must have Read/Write access privileges to change the data on this page.

To access this page, click **Switching > MLD Snooping Querier > MLD Snooping Querier Configuration** in the navigation tree.

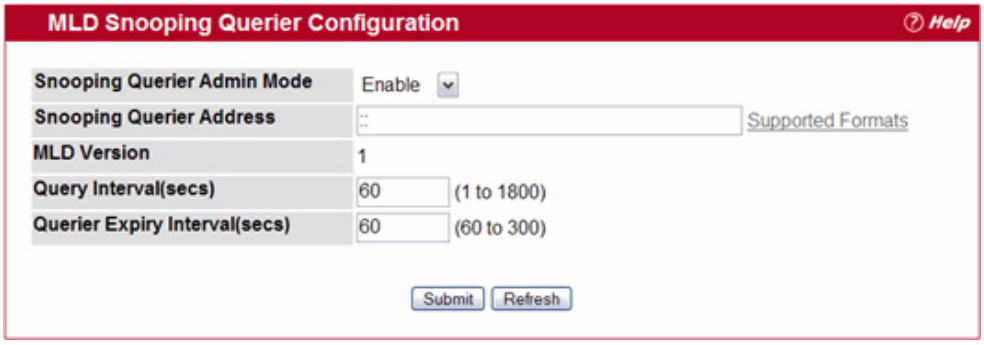

Figure 3-62: MLD Snooping Querier Configuration

Table 3-60: MLD Snooping Querier Configuration Fields

Field	Description
Snooping Querier Admin Mode	Select the administrative mode for MLD Snooping for the switch from the pulldown menu. The default is Disable.
Snooping Querier Address	Specify the Snooping Querier Address to be used as source address in periodic MLD queries. This address is used when no address is configured on the VLAN on which query is being sent.
MLD Version	Specify the MLD protocol version used in periodic MLD queries.
Query Interval	Specify the time interval in seconds between periodic queries sent by the snooping querier. The Query Interval must be a value in the range of 1 and 1800. The default value is 60.
Querier Expiry Interval	Specify the time interval in seconds after which the last querier information is removed. The Querier Expiry Interval must be a value in the range of 60 and 300. The default value is 60.

- If you configure an MLD Snooping querier, click **Submit** to apply the new settings to the switch.
- Click **Refresh** to redisplay the page with the latest information from the switch.

3.15.2 MLD Snooping Querier VLAN Configuration

Use this page to configure MLD queriers for use with VLANs on the network.

To access this page, click **Switching > MLD Snooping Querier > MLD Snooping Querier VLAN Configuration** in the navigation tree.

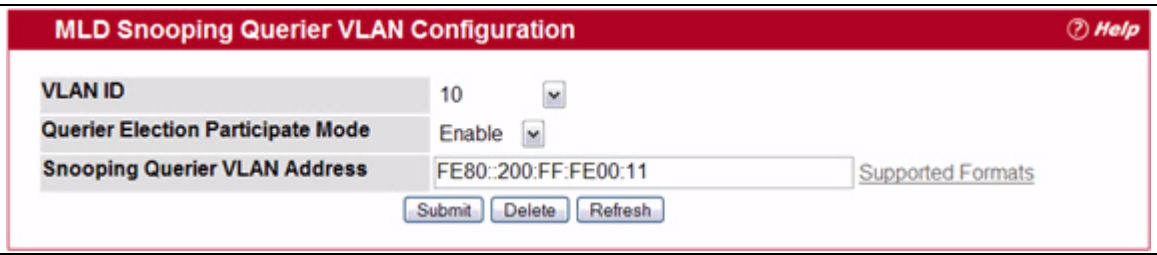

Figure 3-63: MLD Snooping Querier VLAN Configuration

Table 3-61: MLD Snooping Querier VLAN Configuration Fields

Field	Description
VLAN ID	Specifies VLAN ID for which MLD Snooping Querier is to be enabled. You can select New Entry to create a new VLAN ID for the MLD Snooping feature.
Querier Election Participate Mode	Enables or disables Querier Participate Mode. When this mode is disabled, upon seeing another querier of same version in the VLAN, the snooping querier moves to non-querier state. When enabled, the snooping querier participates in querier election, in which the least IP address operates as the querier in that VLAN. The other querier moves to non-querier state.
Snooping Querier VLAN Address	Specifies the Snooping Querier Address to be used as source address in periodic IGMP queries sent on the specified VLAN.

- If you configure or modify the participate mode of a snooping querier for a VLAN, click **Submit** to apply the new settings.
- Click **Refresh** to redisplay the page with the latest information from the switch.
- To remove a querier from the network, select its VLAN ID and click **Delete**.

3.15.3 MLD Snooping Querier VLAN Configuration Summary

Use this page to view summary information for MLD Snooping queriers for on VLANs in the network.

To access this page, click **Switching > MLD Snooping Querier > MLD Snooping Querier VLAN Configuration Summary** in the navigation tree.

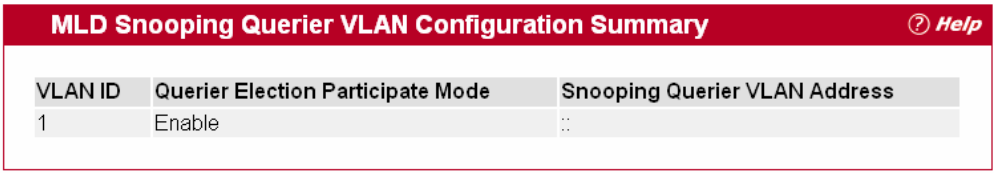

Figure 3-64: MLD Snooping Querier VLAN Configuration Summary

Table 3-62: MLD Snooping Querier VLAN Configuration Summary Fields

Field	Description
VLAN ID	Specifies the VLAN ID on which MLD Snooping Querier is administratively enabled.
Querier Election Participate Mode	Displays the querier election participate mode on the VLAN. When this mode is disabled, up on seeing a query of the same version in the VLAN, the snooping querier moves to non-querier state. When this mode is enabled, the snooping querier participate in querier election, in which the lowest IP address operates as the querier in that VLAN. The other querier moves to non-querier state.
Snooping Querier VLAN Address	Displays the Snooping Querier Address to be used as source address in periodic IGMP queries sent on the specified VLAN.

Click **Refresh** to redisplay the page with the latest information from the router.

3.15.4 MLD Snooping Querier VLAN Status

Use this page to view the operational state and other information for MLD Snooping queriers for VLANs on the network.

To access this page, click **Switching > MLD Snooping Querier > MLD Snooping Querier VLAN Status** in the navigation tree.

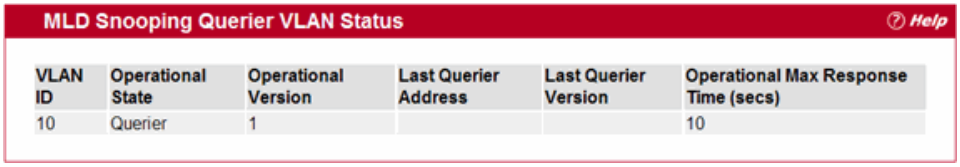

Figure 3-65: MLD Snooping Querier VLAN Status

Table 3-63: MLD Snooping Querier VLAN Status Fields

Field	Description
VLAN ID	Specifies the VLAN ID on which the MLD Snooping Querier is administratively enabled and for which VLAN exists in the VLAN database.
Operational State	Specifies the operational state of the MLD Snooping Querier on a VLAN: • **Querier**: The snooping switch is the querier in the VLAN. The snooping switch will send out periodic queries with a time interval equal to the configured querier query interval. If the snooping switch sees a better querier in the VLAN (i.e., with a numerically lower value), it moves to non-querier mode. • **Non-Querier**: The snooping switch is in non-querier mode in the VLAN. If the querier expiry interval timer expires, the snooping switch moves into querier mode. • **Disabled**: The snooping querier is not operational on the VLAN. The snooping querier moves to disabled mode when MLD Snooping is not operational on the VLAN, when the querier address is not configured, or the network management address is not configured.
Operational Version	Displays the MLD protocol version of the operational querier.
Last Querier Address	Displays the IP address of the last querier from which a query was snooped on the VLAN.
Last Querier Version	Displays the MLD protocol version of the last querier from which a query was snooped on the VLAN.
Operational Max Response Time	Displays the maximum response time to be used in the queries that are sent by the snooping querier.

Click **Refresh** to redisplay the page with the latest information from the switch.

3.16 Creating Port Channels

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

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

Note...

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

Static LAGs are supported. When a port is added to a LAG as a static member, it neither transmits nor receives LACPDUs.

3.16.1 Port Channel Configuration

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

To access the Port Channel Configuration page, click **Switching > Port Channel > Configuration** in the navigation tree.

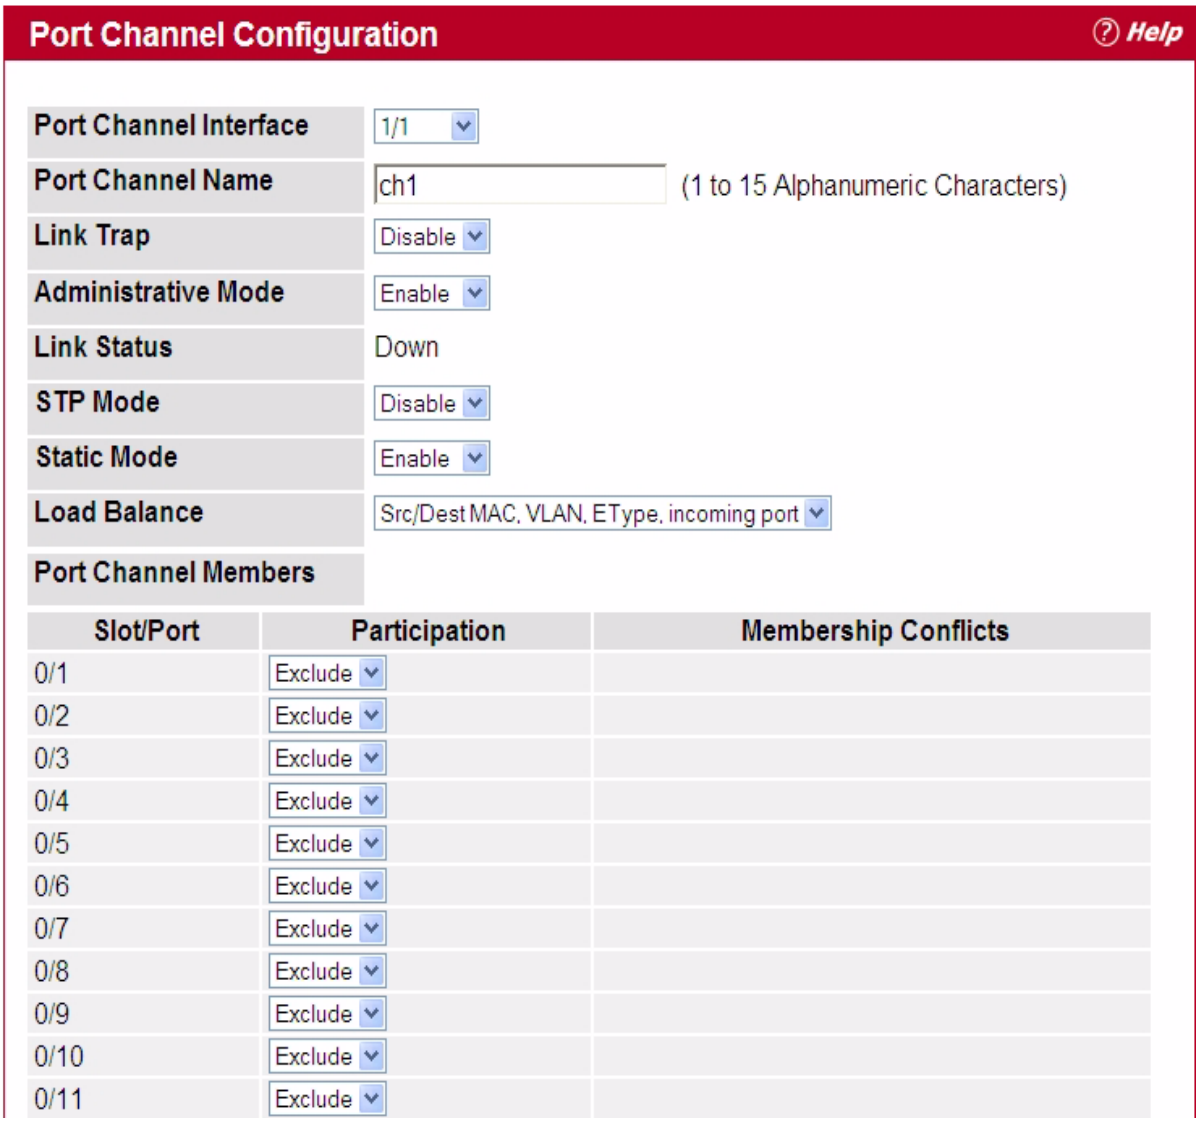

Figure 3-66: Port Channel Configuration

Table 3-64: Port Channel Configuration Fields

Field	Description
Port Channel Name	Select Create from the drop-down menu to configure a new port channel, or select an existing port channel, identified by the interface and name, to modify its settings. The maximum number of port channels is platform-dependent.
Slot/Port	After you create the port channel, this field identifies the Port Channel with the Slot/Port interface naming convention. This field does not appear while you initially configure a new Port Channel.
Port Channel Name	Enter the name you want assigned to the Port Channel. You may enter any string of up to 15 alphanumeric characters. You must specify a valid name in order to create the Port Channel.
Link Trap	Specify whether you want to have a trap sent when link status changes. The factory default is enable, which will cause the trap to be sent.

Table 3-64: Port Channel Configuration Fields (Continued)

Field	Description
Administrative Mode	Select enable or disable from the pulldown menu. When the Port Channel is disabled no traffic will flow and LACPDUs will be dropped, but the links that form the Port Channel will not be released. The factory default is enable.
Link Status	Indicates whether the link is Up or Down.
STP Mode	Select the Spanning Tree Protocol (STP) Administrative Mode associated with the Port Channel: • **Disable**: Spanning tree is disabled for this Port Channel. • **Enable**: Spanning tree is enabled for this Port Channel.
Static Mode	Select enable or disable from the pulldown menu. The factory default is Disable. • **Enable**: The port channel is statically maintained, which means it does not transmit or process received LAGPDUs. The member ports do not transmit LAGPDUs and all the LAGPDUs it may receive are dropped. A static port-channel interface does not require a partner system to be able to aggregate its member ports. • **Disable**: The port channel is dynamically maintained. The interface transmits and processes LAGPDUs and requires a partner system
Load Balance	Select the hashing algorithm used to distribute the traffic load among available physical ports in the LAG. The range of possible values may vary with the type of switch. The possible values are: • Source MAC, VLAN, EtherType, and source port • Destination MAC, VLAN, EtherType and source port • Source/Destination MAC, VLAN, EtherType, and source port • Source IP and Source TCP/UDP Port • Destination IP and Destination TCP/UDP Port • Source/Destination IP and source/destination TCP/UDP Port
Port Channel Members	After you create one or more port channel, this field lists the members of the Port Channel in Slot/Port form. If there are no port channels on the system, this field is not present.
Slot/Port	This column lists the physical ports available on the system.
Participation	Select each port's membership status for the Port Channel you are configuring. There can be a maximum of 8 ports assigned to a Port Channel. • **Include**: The port participates in the port channel. • **Exclude**: The port does not participate in the port channel, which is the default.
Membership Conflicts	Shows ports that are already members of other Port Channels. A port may only be a member of one Port Channel at a time. If the entry is blank, the port is not currently a member of any Port Channel

- If you make any changes to this page, click **Submit** to apply the changes to the system.
- To remove a port channel, select it from the **Port Channel Name** drop-down menu and click delete. All ports that were members of this Port Channel are removed from the Port Channel and included in the default VLAN. This field will not appear when a new Port Channel is being created.

3.16.2 Port Channel Status

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

To access the Port Channel Status page, click **Switching > Port Channel > Status** in the navigation tree.

Port Channel Status Help

Port Channel	Port Channel Name	Port Channel Type	Admin Mode	Link State	STP Mode	Static Mode	Link Trap	Port Channel Members	Active Ports	Load Balance
1/1	ch1	Static	Enable	Down	Disable	Enable	Disable			Src/Dest MAC, VLAN, EType, incoming port
1/2	ch2	Static	Enable	Down	Disable	Enable	Disable			Src/Dest MAC, VLAN, EType, incoming port
1/3	ch3	Static	Enable	Down	Disable	Enable	Disable			Src/Dest MAC, VLAN, EType, incoming port

Figure 3-67: Port Channel Status

Table 3-65: Port Channel Status Fields

Field	Description
Port Channel	Identifies the port channel with the Slot/Port interface naming convention.
Port Channel Name	Identifies the user-configured text name of the port channel.
Port Channel Type	The type of this Port Channel, which is one of the following: • **Static**: The port channel is statically maintained. • **Dynamic**: The port channel is dynamically maintained.
Admin Mode	Select enable or disable from the pulldown menu. When the Port Channel is disabled no traffic will flow and LACPDUs will be dropped, but the links that form the Port Channel will not be released. The factory default is enable.
Link State	Indicates whether the link is Up or Down.
STP Mode	Shows whether the Spanning Tree Protocol (STP) Administrative Mode is enabled or disabled on the port channel
Link Trap	Shows whether to send traps when link status changes. If the status is Enabled, traps are sent.
Static Mode	Shows whether static mode is enabled for this port channel.
Configured Ports	Lists the ports that are members of the Port Channel, in Slot/Port notation. There can be a maximum of 8 ports assigned to a Port Channel.
Active Ports	Lists the ports that are actively participating members of this Port Channel, in Slot/Port notation.
Load Balance	Shows the hashing algorithm used to distribute the traffic load among available physical ports in the LAG. The range of possible values may vary with the type of switch. The possible values are: • Source MAC, VLAN, EtherType, and source port • Destination MAC, VLAN, EtherType and source port • Source/Destination MAC, VLAN, EtherType, and source port • Source IP and Source TCP/UDP Port • Destination IP and Destination TCP/UDP Port • Source/Destination IP and source/destination TCP/UDP Port

3.17 Viewing Multicast Forwarding Database Information

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

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

This Multicast Support folder contains links to the following pages:

- MFDB Table
- MFDB GMRP Table
- MFDB IGMP Snooping Table
- MFDB MLD Snooping Table
- MFDB Statistics

3.17.1 MFDB Table

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

To access the MFDB Table page, click **Switching > Forwarding Database > MFDB Table** in the navigation tree.

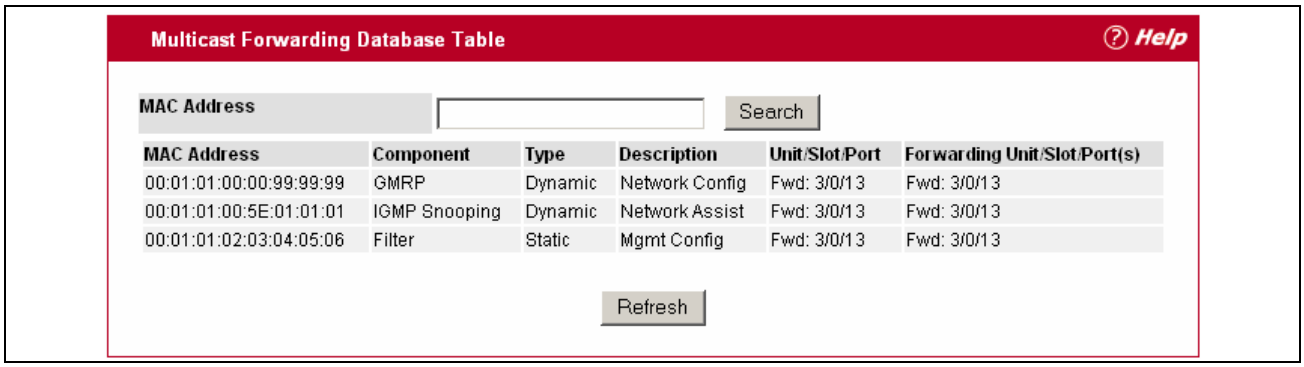

Figure 3-68: MFDB Table

Table 3-66: MFDB Table Fields

Field	Description
MAC Address	Enter the VLAN ID/MAC Address pair whose MFDB table entry you want displayed. Enter eight two-digit hexadecimal numbers separated by colons, for example 00:01:23:43:45:67:89:AB. The first two 2-digit hexadecimal numbers are the VLAN ID and the remaining numbers are the MAC address. Click on the "Search" button. If the address exists, that entry will be displayed. An exact match is required.
MAC Address	The multicast MAC address for which you requested data.
Component	This is the component that is responsible for this entry in the Multicast Forwarding Database. Possible values are MLD Snooping, GMRP,IGMP Snooping, and Static Filtering.
Type	This displays the type of the entry. Static entries are those that are configured by the end user. Dynamic entries are added to the table as a result of a learning process or protocol.
Description	The text description of this multicast table entry. Possible values are Management Configured, Network Configured and Network Assisted.
Slot/Port	The list of interfaces that are designated for forwarding (Fwd) and filtering (Flt) for the selected address.
Forwarding Slot/Port(s)	The resultant forwarding list is derived from combining all the forwarding interfaces and removing the interfaces that are listed as the static filtering interfaces.

- To search for a MAC address if the list is too long to scan, enter the MAC address in hex format and click **Search**.
- Click **Refresh** to update the information on the screen with the most current data.

3.17.2 MFDB GMRP Table

Use the GMRP Table page to view all of the entries in the Multicast Forwarding Database that were created for the GARP Multicast Registration Protocol.

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

Figure 3-69: GMRP Table

Table 3-67: GMRP Table Fields

Field	Description
MAC Address	A VLAN ID/multicast MAC address pair for which the switch has forwarding and or filtering information. The format is 8 two-digit hexadecimal numbers that are separated by colons, for example 00:01:23:45:67:89:AB:CD.
Type	This displays the type of the entry. Static entries are those that are configured by the end user. Dynamic entries are added to the table as a result of a learning process or protocol.
Description	The text description of this multicast table entry. Possible values are Management Configured, Network Configured and Network Assisted.
Slot/Port	The list of interfaces that are designated for forwarding (Fwd) and filtering (Flt) for the associated address.

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

3.17.3 MFDB IGMP Snooping Table

Use the IGMP Snooping Table page to view all of the entries in the Multicast Forwarding Database that were created for IGMP snooping.

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

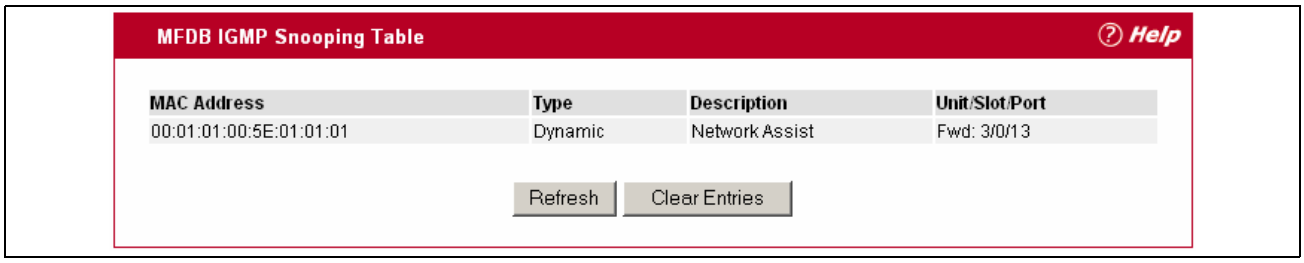

Figure 3-70: IGMP Snooping Table

Table 3-68: MFDB IGMP Snooping Table Fields

Field	Description
MAC Address	A VLAN ID/multicast MAC address pair for which the switch has forwarding and or filtering information. The format is 8 two-digit hexadecimal numbers that are separated by colons, for example, 00:01:23:45:67:89:AB:CD.
Type	This displays the type of the entry. Static entries are those that are configured by the end user. Dynamic entries are added to the table as a result of a learning process or protocol.
Description	The text description of this multicast table entry. Possible values are Management Configured, Network Configured and Network Assisted.
Slot/Port	The list of interfaces that are designated for forwarding (Fwd) and filtering (Flt) for the associated address.

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

- Click **Clear Entries** to tell the IGMP Snooping component to delete all of its entries from the multicast forwarding database.

3.17.4 MFDB MLD Snooping Table

Use the MLD Snooping Table page to view all of the entries in the Multicast Forwarding Database that were created for MLD Snooping.

To access the MLD Snooping Table page, click **Switching > Multicast Forwarding Database > MLD Snooping Table** in the navigation tree.

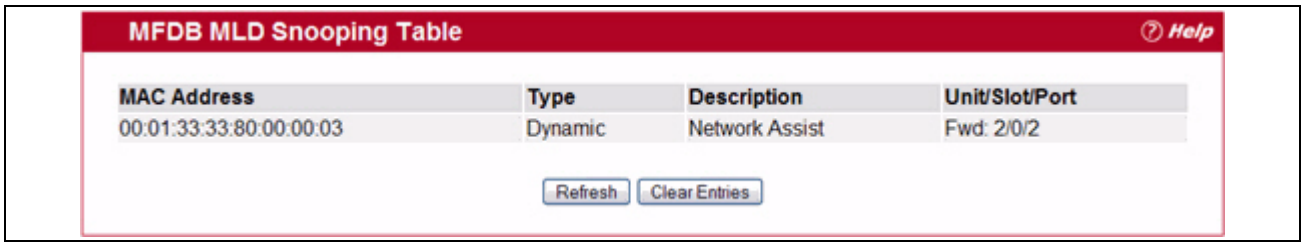

Figure 3-71: MLD Snooping Table

Table 3-69: MLD Snooping Table Fields

Field	Description
MAC Address	A VLAN ID/multicast MAC address pair for which the switch has forwarding and or filtering information. The format is 8 two-digit hexadecimal numbers that are separated by colons, for example, 00:01:23:45:67:89:AB:CD.
Type	This displays the type of the entry. Static entries are those that are configured by the end user. Dynamic entries are added to the table as a result of a learning process or protocol.
Description	The text description of this multicast table entry. Possible values are Management Configured, Network Configured and Network Assisted.
Slot/Port	The list of interfaces that are designated for forwarding (Fwd) and filtering (Flt) for the associated address.

- Click **Refresh** to update the information on the screen with the most current data.
- Click **Clear Entries** to tell the MLD Snooping component to delete all of its entries from the multicast forwarding database.

3.17.5 MFDB Statistics

Use the multicast forwarding database Stats page to view statistical information about the MFDB table.

To access the Stats page, click **Switching > Multicast Forwarding Database > Stats** in the navigation tree.

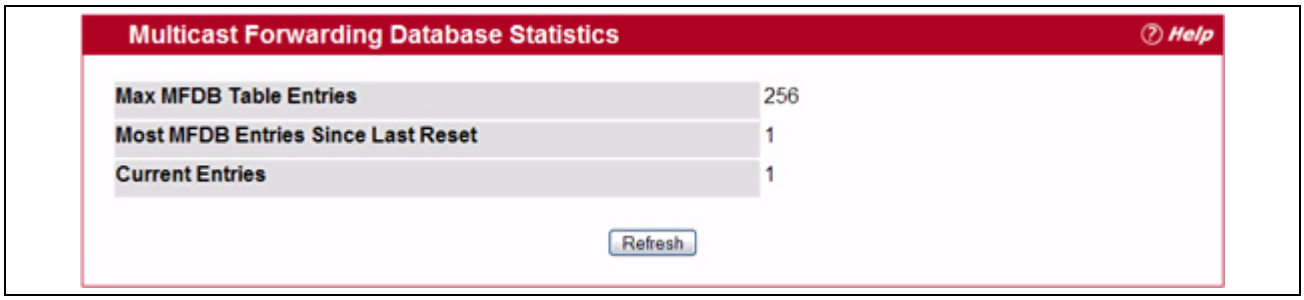

Figure 3-72: Multicast Forwarding Database Statistics

Table 3-70: Multicast Forwarding Database Statistics Fields

Field	Description
Max MFDB Entries	Shows the maximum number of entries that the Multicast Forwarding Database table can hold.
Most MFDB Entries Since Last Reset	The largest number of entries that have been present in the Multicast Forwarding Database table since the system was last reset. This value is also known as the MFDB high-water mark.
Current Entries	Shows the current number of entries in the Multicast Forwarding Database table.

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

3.18 Configuring Spanning Tree Protocol

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

Classic STP provides a single path between end stations, avoiding and eliminating loops. For information on configuring Common STP, see 3.18.4 CST Port Configuration/Status.

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

Note...

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

The Spanning Tree folder contains links to the following STP pages:

- Switch Configuration/Status
- CST Configuration/Status
- MST Configuration/Status
- CST Port Configuration/Status
- MST Port Configuration/Status
- Statistics

3.18.1 Switch Configuration/Status

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

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

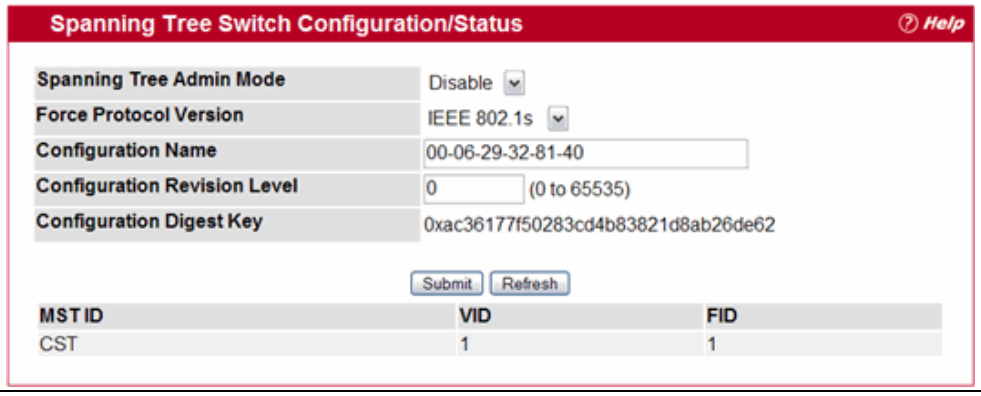

Figure 3-73: Spanning Tree Switch Configuration/Status

Table 3-71: Spanning Tree Switch Configuration/Status Fields

Field	Description
Spanning Tree Admin Mode	Enables or disables STP on the switch.
Force Protocol Version	Specifies the Force Protocol Version parameter for the switch: • **IEEE 802.1D**: Spanning Tree Protocol (STP) • **IEEE 802.1w**: Rapid Spanning Tree Protocol (RSTP) • **IEEE 802.1s**: Multiple Spanning Tree Protocol (MSTP)
Configuration Name	Name used to identify the configuration currently being used. It may be up to 32 alphanumeric characters.
Configuration Revision Level	Number used to identify the configuration currently being used. The values allowed are between 0 and 65535. The default value is 0.

Table 3-71: Spanning Tree Switch Configuration/Status Fields (Continued)

Field	Description
Configuration Digest Key	Number used to identify the configuration currently being used. The digest key is generated based on the association of VLANs to different instances. To ensure the digest key is same on two different switches, the mapping of VLAN-to-instance must be the same.
MST ID	Table consisting of the MST instances (including the CST) and the corresponding VLAN IDs associated with each of them.
VID	This table consists of the VLAN identifier (VID) and the corresponding filtering identifier (FID) associated with each VID.
FID	Table consisting of the FIDs and the corresponding VLAN IDs associated with each of them.

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

3.18.2 CST Configuration/Status

Use the Spanning Tree CST Configuration/Status page to configure Common Spanning Tree (CST) and Internal Spanning Tree on the switch.

To display the Spanning Tree CST Configuration/Status page, click **Switching** > **Spanning Tree** > **CST Configuration/Status** in the navigation tree.

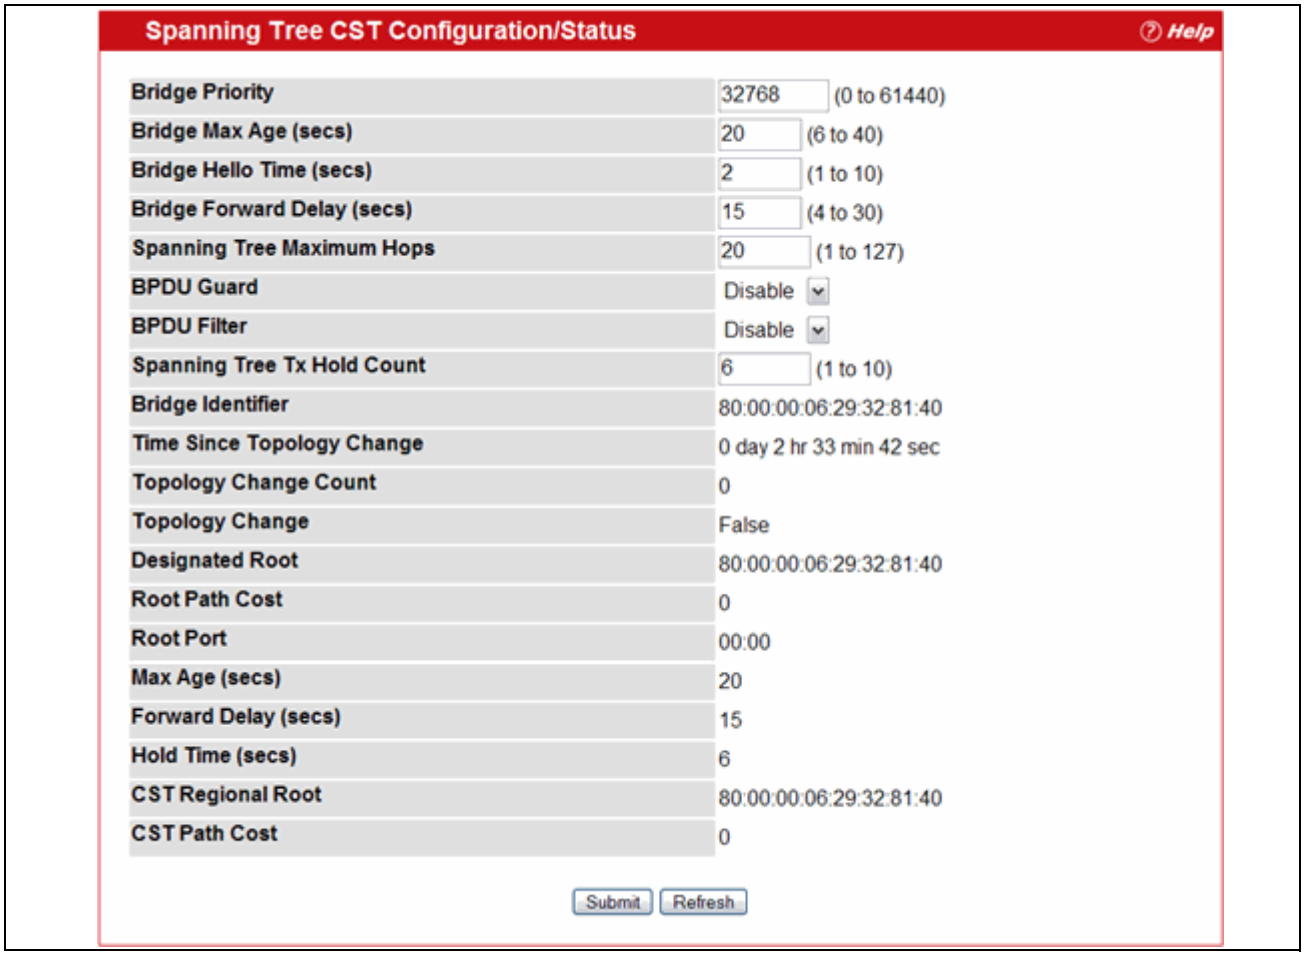

Figure 3-74: Spanning Tree CST Configuration/Status

Table 3-72: Spanning Tree CST Configuration/Status Fields

Field	Description
Bridge Priority	Specifies the bridge priority value. When switches or bridges are running STP, each is assigned a priority. After exchanging BPDUs, the switch with the lowest priority value becomes the root bridge. The bridge priority is a multiple of 4096. If you specify a priority that is not a multiple of 4096, the priority is automatically set to the next lowest priority that is a multiple of 4096. For example if the priority is attempted to be set to any value between 0 and 4095, it will be set to 0. The default priority is 32768.The valid range is 0-61440.
Bridge Max Age (secs)	Specifies the switch maximum age time, which indicates the amount of time in seconds a bridge waits before implementing a topological change. The valid range is 6-40, and the value must be less than or equal to (2 * Bridge Forward Delay) – 1 and greater than or equal to 2 * (Bridge Hello Time +1). The default value is 20.
Bridge Hello Time (secs)	Specifies the switch Hello time, which indicates the amount of time in seconds a root bridge waits between configuration messages. The valid range is 1-10, and the default value is 2. The value must be less than or equal to (Bridge Max Age / 2) – 1. The default hello time value is 2.
Bridge Forward Delay (secs)	Specifies the switch forward delay time, which indicates the amount of time in seconds a bridge remains in a listening and learning state before forwarding packets. The value must be greater or equal to (Bridge Max Age / 2) + 1. The time range is from 4 seconds to 30 seconds. The default value is 15.
Spanning Tree Maximum Hops	Specifies the maximum number of bridge hops the information for a particular CST instance can travel before being discarded.
BPDU Guard	Enable or disable the BPDU Guard.The switches behind the edge ports that have BPDU guard enabled will not be able to influence the overall STP topology. Using the BPDU Guard feature can help enforce the STP domain borders and keep the active topology be consistent and predictable.
BPUD Filter	Enable or disable the BPDU Filter. When BPDU filtering is enabled, the port drops the BPDUs received.
Spanning Tree Tx Hold Count	Configure the maximum number of BPDUs the bridge is allowed to send within the hello time window.The default value is 6.
Bridge Identifier	The bridge identifier for the CST. It is made up using the bridge priority and the base MAC address of the bridge.
Time Since Topology Change	Displays the total amount of time since the last topographic change. The time is displayed in hour/minute/second format, for example, 5 hours 10 minutes and 4 seconds.
Topology Changes Counts	Displays the total amount of STP state changes that have occurred.
Topology Change	Indicates whether a topology change is in progress on any port assigned to the CST. The possible values are True or False.
Designated Root	Displays the bridge identifier of the root bridge, which is made up from the bridge priority and the base MAC address of the bridge.
Root Path Cost	Displays the cost of the path from this bridge to the designated root.
Root Port	Indicates the root port of the selected instance.
Max Age	Shows the path Cost to the Designated Root for the CST.
Forward Delay	Shows the derived value of the Root Port Bridge Forward Delay parameter.
Hold Time	Indicates the minimum time between transmission of Configuration BPDUs.
CST Regional Root	Shows the priority and base MAC address of the CST Regional Root.
CST Path Cost	Shows the path Cost to the CST tree Regional Root.

3.18.3 MST Configuration/Status

Use the Spanning Tree MST Configuration/Status page to configure Multiple Spanning Tree (MST) on the switch.

To display the Spanning Tree MST Configuration/Status page, click **Switching** > **Spanning Tree** > **MST Configuration/Status** in the navigation tree.

If no MST instances exist, or if you select Create from the **MST** field, the MST Configuration/Status page looks like the screen in Figure 3-75.

Figure 3-75: Spanning Tree MST Configuration/Status

Figure 3-76 shows an example of the page with an MST instance configured.

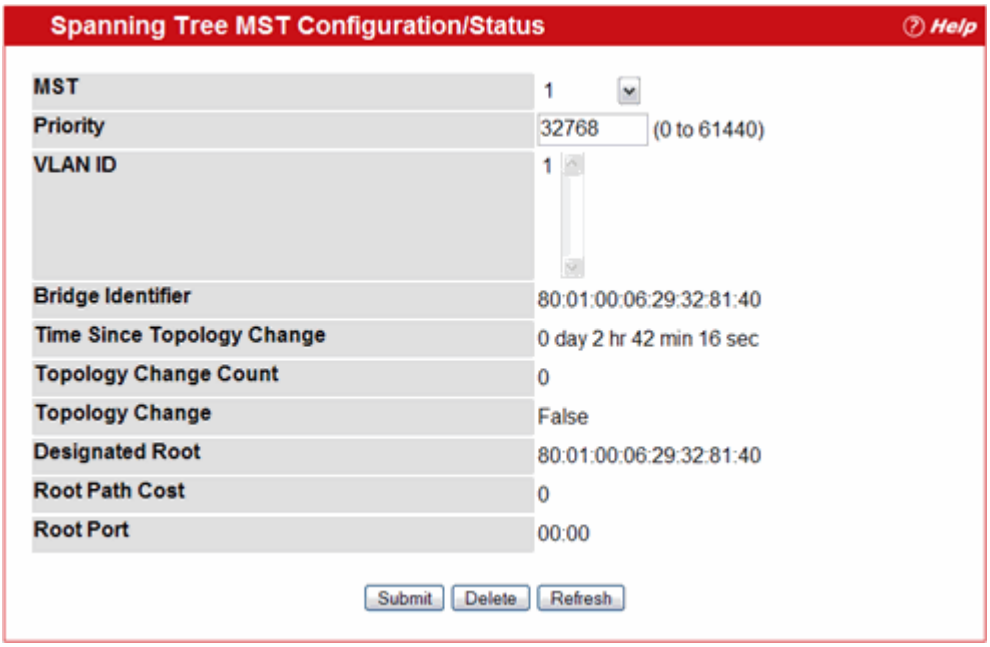

Figure 3-76: Spanning Tree MST Configuration/Status

Table 3-73: Spanning Tree MST Configuration/Status

Field	Description
MST	Use the drop-down menu to create and configure a new MST or select an existing MST to display or configure.
MST ID	This is only visible when Create is selected from the **MST** field drop-down menu. The ID of the MST being created. Valid values for this are between 1 and 4094.
Priority	Specifies the bridge priority value for the MST. When switches or bridges are running STP, each is assigned a priority. After exchanging BPDUs, the switch with the lowest priority value becomes the root bridge. The bridge priority is a multiple of 4096. If you specify a priority that is not a multiple of 4096, the priority is automatically set to the next lowest priority that is a multiple of 4096. For example if the priority is attempted to be set to any value between 0 and 4095, it will be set to 0. The default priority is 32768.The valid range is 0-61440.
VLAN ID	This gives a list box of all VLANs on the switch. The VLANs associated with the MST instance which is selected are highlighted on the list. These can be selected or unselected for reconfiguring the association of VLANs to MST instances.
Bridge Identifier	The bridge identifier for the selected MST instance. It is made up using the bridge priority and the base MAC address of the bridge.
Time Since Topology Change	Displays the total amount of time since the last topographic change. The time is displayed in hour/minute/second format, for example, 5 hours 10 minutes and 4 seconds.
Topology Changes Counts	Displays the total number of MST state changes that have occurred.
Topology Change	Indicates whether a topology change is in progress on any port assigned to the CST. The possible values are True or False.
Designated Root	Displays the bridge identifier of the root bridge, which is made up from the bridge priority and the base MAC address of the bridge.
Root Path Cost	Displays the path cost to the Designated Root for this MST instance.
Root Port	Indicates the port to access the Designated Root for this MST instance.

- If you make any configuration changes, click Submit to apply the new settings to the switch.
- Click **Force** to force the port to send out 802.1w or 802.1D BPDUs.
- Click **Refresh** to update the screen with most recent data.

3.18.4 CST Port Configuration/Status

Use the Spanning Tree CST Port Configuration/Status page to configure Common Spanning Tree (CST) and Internal Spanning Tree on a specific port on the switch.

To display the Spanning Tree CST Port Configuration/Status page, click **Switching** > **Spanning Tree** > **CST Port Configuration/Status** in the navigation tree.

Spanning Tree CST Configuration/Status Help

Bridge Priority	32768	(0 to 61440)
Bridge Max Age (secs)	20	(6 to 40)
Bridge Hello Time (secs)	2	
Bridge Forward Delay (secs)	15	(4 to 30)
Spanning Tree Maximum Hops	20	(6 to 40)
BPDU Guard	Disable	
BPDU Filter	Disable	
Spanning Tree Tx Hold Count	6	(1 to 10)
Bridge Identifier	80:00:00:a0:a5:5d:2b:81	
Time Since Topology Change	0 day 3 hr 37 min 53 sec	
Topology Change Count	0	
Topology Change	False	
Designated Root	80:00:00:a0:a5:5d:2b:81	
Root Path Cost	0	
Root Port	00:00	
Max Age (secs)	20	
Forward Delay (secs)	15	
Hold Time (secs)	6	
CST Regional Root	80:00:00:a0:a5:5d:2b:81	
CST Path Cost	0	

Submit Refresh

Figure 3-77: Spanning Tree CST Port Configuration/Status

Table 3-74: Spanning Tree CST Port Configuration/Status Fields

Field	Description
Slot/Port	Select a physical or port channel interface to configure. The port is associated with the VLAN(s) associated with the CST.
Port Priority	The priority for a particular port within the CST. The port priority is set in multiples of 16. If you specify a value that is not a multiple of 16, the priority is set to the priority is automatically set to the next lowest priority that is a multiple of 16. For example, if you set a value between 0 and 15, the priority is set to 0. If you specify a number between 16 and 31, the priority is set to 16.
Admin Edge Port	Determines whether the specified port is an Edge Port within the CIST. It takes a value of TRUE or FALSE, where the default value is FALSE.
Port Path Cost	Set the Path Cost to a new value for the specified port in the common and internal spanning tree. It takes a value in the range of 1 to 200000000.
Auto-calculate Port Path Cost	Displays whether the path cost is automatically calculated (Enabled) or not (Disabled). Path cost is calculated based on the link speed of the port if the configured value for Port Path Cost is zero.
Hello Timer	Specifies the switch Hello time, which indicates the amount of time in seconds a port waits between configuration messages. The valid range is 1-10, and the default value is 2. The value must be less than or equal to (Bridge Max Age / 2) – 1. The default hello time value is 2.
External Port Path Cost	Set the External Path Cost to a new value for the specified port in the spanning tree. It takes a value in the range of 1 to 200000000.
Auto-calculate External Port Path Cost	Displays whether the external path cost is automatically calculated (Enabled) or not (Disabled). External Path cost will be calculated based on the link speed of the port if the configured value for External Port Path Cost is zero.
BPDU Filter	Enable or disable the BPDU Filter, which filters the BPDU traffic on this port when STP is enabled on this port.
BPDU Flood	Enable or disable the BPDU Flood, which floods the BPDU traffic arriving on this port when STP is disabled on this port.
BPDU Guard Effect	If BPDU Guard is enabled for the switch and the edge port receives a BPDU, the port will be disabled and the status of this field is Enabled.
Port ID	The port identifier for the specified port within the CST. It is made up from the port priority and the interface number of the port.
Port Up Time Since Counters Last Cleared	Time since the counters were last cleared, displayed in Days, Hours, Minutes, and Seconds.
Port Mode	Spanning Tree Protocol Administrative Mode associated with the port or port channel. The possible values are Enable or Disable.
Port Forwarding State	Indicates the current STP state of a port. If enabled, the port state determines what forwarding action is taken on traffic. Possible port states are: • **Disabled**: STP is currently disabled on the port. The port forwards traffic while learning MAC addresses. • **Blocking**: The port is currently blocked and cannot be used to forward traffic or learn MAC addresses. • **Listening**: The port is currently in the listening mode. The port cannot forward traffic nor can it learn MAC addresses. • **Learning**: The port is currently in the learning mode. The port cannot forward traffic, however, it can learn new MAC addresses. • **Forwarding**: The port is currently in the forwarding mode. The port can forward traffic and learn new MAC addresses.
Port Role	Each MST Bridge Port that is enabled is assigned a Port Role for each spanning tree. The port role will be one of the following values: Root Port, Designated Port, Alternate Port, Backup Port, Master Port or Disabled Port.
Designated Root	Root Bridge for the CST. It is made up using the bridge priority and the base MAC address of the bridge.

Table 3-74: Spanning Tree CST Port Configuration/Status Fields (Continued)

Field	Description
Designated Cost	Displays cost of the port participating in the STP topology. Ports with a lower cost are less likely to be blocked if STP detects loops.
Designated Bridge	Bridge Identifier of the bridge with the Designated Port. It is made up using the bridge priority and the base MAC address of the bridge.
Designated Port	Port Identifier on the Designated Bridge that offers the lowest cost to the LAN. It is made up from the port priority and the interface number of the port.
Topology Change Acknowledge	Identifies whether the next BPDU to be transmitted for this port would have the topology change acknowledgement flag set. It is either "True" or "False".
Auto Edge	Configuring the auto edge mode of a port allows the port to become an edge port if it does not see BPDUs for some duration. The possible values are Enable or Disable.
Edge Port	Indicates whether the port is enabled as an edge port.
Point-to-point MAC	Derived value of the point-to-point status.
Root Guard	Configuring the root guard mode sets a port to discard any superior information received by the port and thus protect against root of the device from changing. The port gets put into discarding state and does not forward any packets.The possible values are Enable or Disable.
Loop Guard	Configuring the loop guard mode prevents a port from erroneously transitioning from blocking state to forwarding when the port stops receiving BPDUs. The port is marked as being in loop-inconsistent state. In this state, the port does not forward packets. The possible values are Enable or Disable.
TCN Guard	Configuring the TCN guard for a port restricts the port from propagating any topology change information received through that port.The possible values are Enable or Disable.
CST Regional Root	Shows the priority and base MAC address of the CST Regional Root.
CST Path Cost	Shows the path Cost to the CST tree Regional Root.
Loop Inconsistent State	Identifies whether the port is currently in a loop inconsistent state. If the port is in a loop inconsistent state, it does not forward packets.
Transitions Into Loop Inconsistent State	Shows the number of times this interface has moved into a loop inconsistent state.
Transitions Out Of Loop Inconsistent State	Shows the number of times this interface has gotten out of a loop inconsistent state.

- If you make any configuration changes, click **Submit** to apply the new settings to the switch.
- Click **Force** to force the port to send out 802.1w or 802.1D BPDUs.
- Click **Refresh** to update the screen with most recent data.

3.18.5 MST Port Configuration/Status

Use the Spanning Tree MST Port Configuration/Status page to configure Multiple Spanning Tree (MST) on a specific port on the switch.

To display the Spanning Tree MST Port Configuration/Status page, click **Switching > Spanning Tree > MST Port Configuration/Status** in the navigation tree.

Note...

If no MST instances have been configured on the switch, the page displays a “No MSTs Available” message and does not display the fields shown in Figure 3-78.

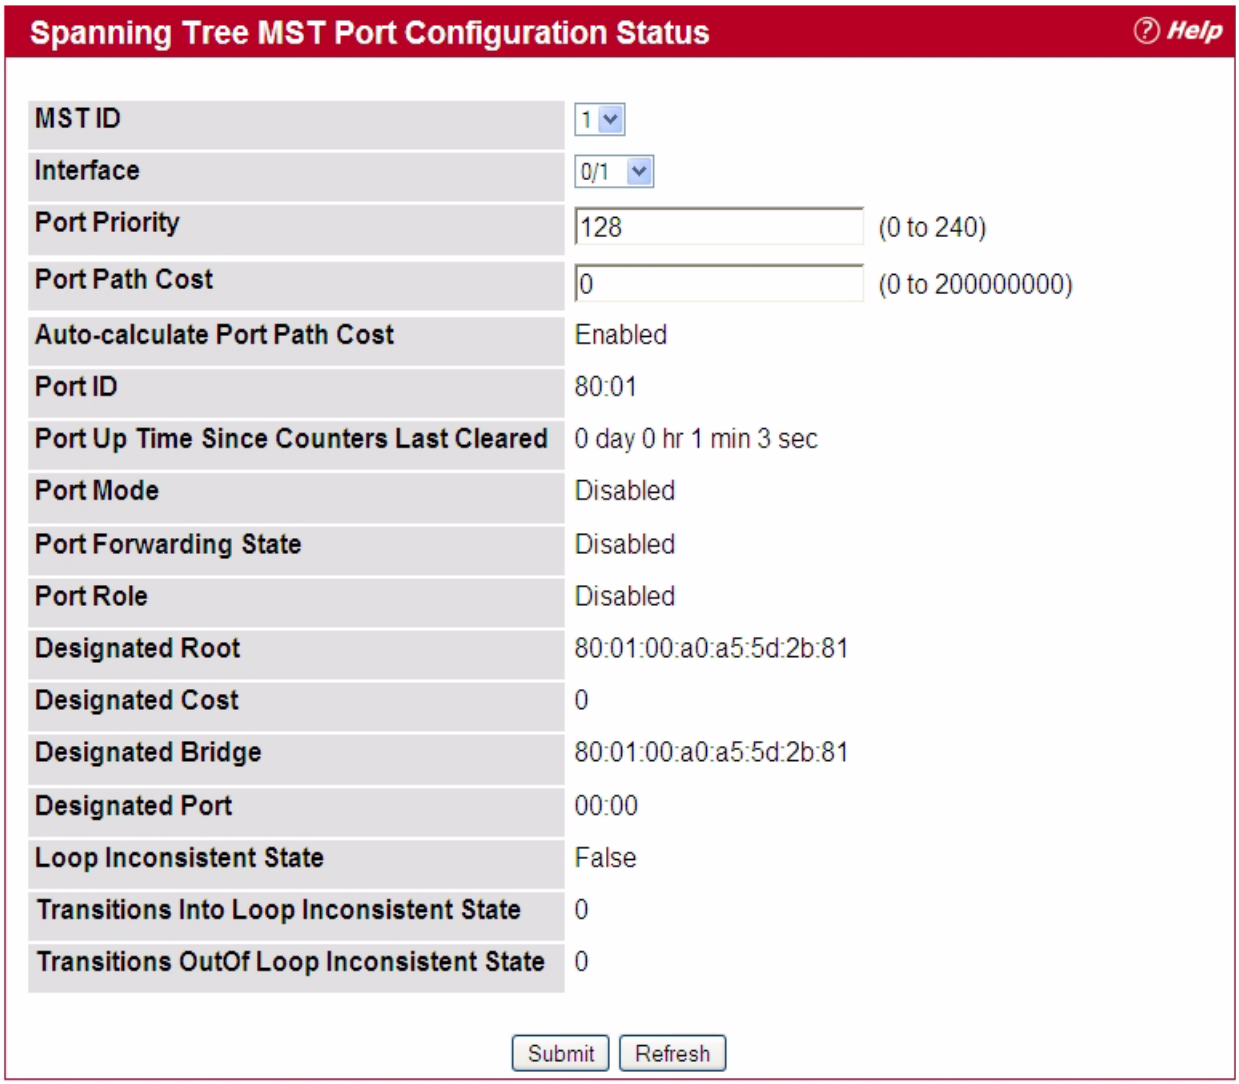

Figure 3-78: Spanning Tree MST Port Configuration/Status

Table 3-75: Spanning Tree MST Port Configuration/Status Fields

Field	Description
MST ID	Select an existing MST instance from drop-down list to display or configure its values.
Slot/Port	Select a physical or port channel interface to configure. The port is associated with the VLAN(s) associated with the MST.
Port Priority	The priority for a particular port within the MST. The port priority is set in multiples of 16. If you specify a value that is not a multiple of 16, the priority is set to the priority is automatically set to the next lowest priority that is a multiple of 16. For example, if you set a value between 0 and 15, the priority is set to 0. If you specify a number between 16 and 31, the priority is set to 16.
Port Path Cost	Set the Path Cost to a new value for the specified port in the selected MST instance. It takes a value in the range of 1 to 200000000.

Table 3-75: Spanning Tree MST Port Configuration/Status Fields (Continued)

Field	Description
Auto-calculate Port Path Cost	Displays whether the path cost is automatically calculated (Enabled) or not (Disabled). Path cost is calculated based on the link speed of the port if the configured value for Port Path Cost is zero.
Port ID	The port identifier for the specified port within the CST. It is made up from the port priority and the interface number of the port.
Port Up Time Since Counters Last Cleared	Time since the counters were last cleared, displayed in Days, Hours, Minutes, and Seconds.
Port Mode	Shows whether STP is enabled on the port. To enable STP on a port, use the **System > Port > Configuration** page.
Port Forwarding State	Indicates the current STP state of a port. If enabled, the port state determines what forwarding action is taken on traffic. Possible port states are: • **Disabled**: STP is currently disabled on the port. The port forwards traffic while learning MAC addresses. • **Blocking**: The port is currently blocked and cannot be used to forward traffic or learn MAC addresses. • **Listening**: The port is currently in the listening mode. The port cannot forward traffic nor can it learn MAC addresses. • **Learning**: The port is currently in the learning mode. The port cannot forward traffic, however, it can learn new MAC addresses. • **Forwarding**: The port is currently in the forwarding mode. The port can forward traffic and learn new MAC addresses
Port Role	Each MST Bridge Port that is enabled is assigned a Port Role for each spanning tree. The port role will be one of the following values: Root Port, Designated Port, Alternate Port, Backup Port, Master Port or Disabled Port.
Designated Root	Root Bridge for the selected MST instance. It is made up using the bridge priority and the base MAC address of the bridge.
Designated Cost	Displays cost of the port participating in the STP topology. Ports with a lower cost are less likely to be blocked if STP detects loops.
Designated Bridge	Bridge Identifier of the bridge with the Designated Port. It is made up using the bridge priority and the base MAC address of the bridge.
Designated Port	Port Identifier on the Designated Bridge that offers the lowest cost to the LAN. It is made up from the port priority and the interface number of the port.
Loop Inconsistent State	This parameter identifies whether the port is in a loop inconsistent state in the specified MST instance. If the port is in a loop inconsistent state, it does not forward packets.
Transitions Into Loop Inconsistent State	Shows the number of times this interface has gone into a loop inconsistent state.
Transitions Out Of Loop Inconsistent State	Shows the number of times this interface has gotten out of a loop inconsistent state.

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

3.18.6 Statistics

Use the Spanning Tree Statistics page to view information about the number and type of bridge protocol data units (BPDUs) transmitted and received on each port.

To display the Spanning Tree Statistics page, click **Switching > Spanning Tree > Statistics** in the navigation tree.

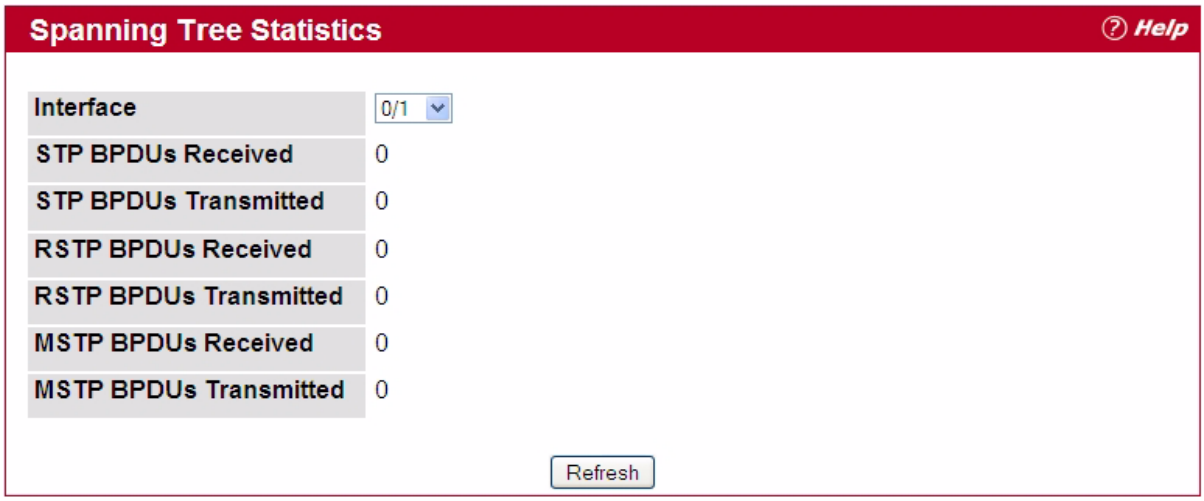

Figure 3-79: Spanning Tree Statistics

Table 3-76: Spanning Tree Statistics Fields

Field	Description
Slot/Port	Select a physical or port channel interface to view its statistics.
STP BPDUs Received	Number of STP BPDUs received at the selected port.
STP BPDUs Transmitted	Number of STP BPDUs transmitted from the selected port.
RSTP BPDUs Received	Number of RSTP BPDUs received at the selected port.
RSTP BPDUs Transmitted	Number of RSTP BPDUs transmitted from the selected port.
MSTP BPDUs Received	Number of MSTP BPDUs received at the selected port.
MSTP BPDUs Transmitted	Number of MSTP BPDUs transmitted from the selected port.

Click **Refresh** to update the screen with most recent data.

3.19 Mapping 802.1p Priority

The IEEE 802.1p feature allows traffic prioritization at the MAC level. The switch can prioritize traffic based on the 802.1p tag attached to the L2 frame. Each port on the switch has multiple queues to give preference to certain packets over others based on the class of service (CoS) criteria you specify. When a packet is queued for transmission in a port, the rate at which it is serviced depends on how the queue is configured and possibly the amount of traffic present in the other queues of the port. If a delay is necessary, packets get held in the queue until the scheduler authorizes the queue for transmission.

Use the 802.1p Priority Mapping page in the Class of Service folder to assign 802.1p priority values to various traffic classes on one or more interfaces.

To display the page, click **Switching > Class of Service** > **802.1p Priority Mapping** in the navigation tree.

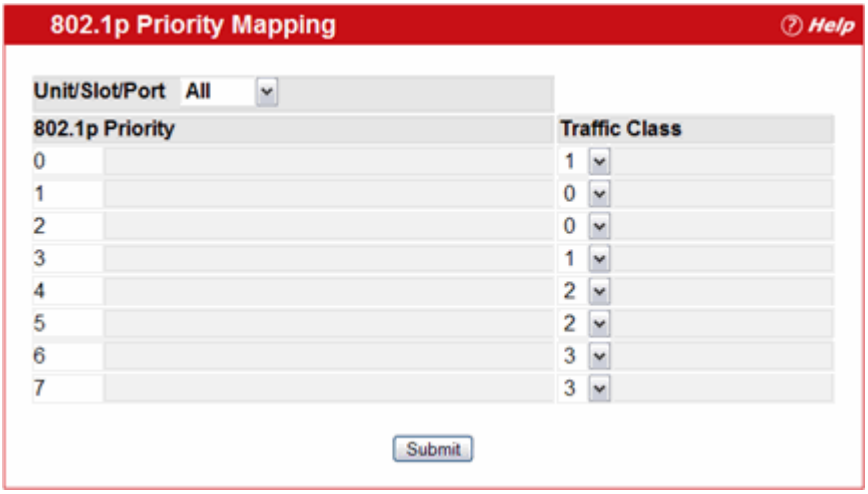

Figure 3-80: 802.1p Priority Mapping

Table 3-77: 802.1p Priority Mapping

Field	Description
Slot/Port	Selects the interface to which the class of service configuration is applied.
802.1p Priority	Displays the 802.1p priority to be mapped. Priority goes from low (0) to high (7). For example, traffic with a priority of 0 is for most data traffic and is sent using "best effort." Traffic with a higher priority, such as 6, might be time-sensitive traffic, such as voice or video.
Traffic Class	The traffic class is the hardware queue for a port. Higher traffic class values indicate a higher queue position. Before traffic in a lower queue is sent, it must wait for traffic in higher queues to be sent. To change the default priority-to-queue mapping, select a new traffic class value from the drop-down menu.

If you make any changes to the page, click **Submit** to apply the new values to the system.

3.20 Configuring Port Security

Port Security can be enabled on a per-port basis. When a port is locked, only packets with allowable source MAC addresses can be forwarded. All other packets are discarded. A MAC address can be defined as allowable by one of two methods: dynamically or statically. Note that both methods are used concurrently when a port is locked.

Dynamic locking implements a "first arrival" mechanism for Port Security. You specify how many addresses can be learned on the locked port. If the limit has not been reached, a packet with an unknown source MAC address is learned and forwarded normally. Once the limit is reached, no more addresses are learned on the port. Any packets with source MAC addresses that were not already learned are discarded. Note that you can effectively disable dynamic locking by setting the number of allowable dynamic entries to zero.

Static locking allows you to specify a list of MAC addresses that are allowed on a port. The behavior of packets is the same as for dynamic locking: only packets with an allowable source MAC address can be forwarded.

To see the MAC addresses learned on a specific port, see Configuring and Searching the Forwarding Database.

Disabled ports can only be activated from the **Configuring Ports** page.

The **Port Security** folder contains links to the following pages:

- Port Security Administration
- Port Security Interface Configuration
- Port Security Static
- Port Security Dynamic
- Port Security Violation Status

3.20.1 Port Security Administration

Use the Port Security Administration page to enable or disable the port security feature on your switch.

To access the Port Security Administration page, click **Switching > Port Security > Port Security Administration** in the navigation tree.

Figure 3-81: Port Security Administration

Select **Enable** or **Disable** from the Port Security Mode list and click **Submit**.

3.20.2 Port Security Interface Configuration

Use this page to configure the port security feature on a selected interface.

To access the Port Security Interface Configuration page, click **Switching > Port Security > Port Security Interface Configuration** in the navigation tree.

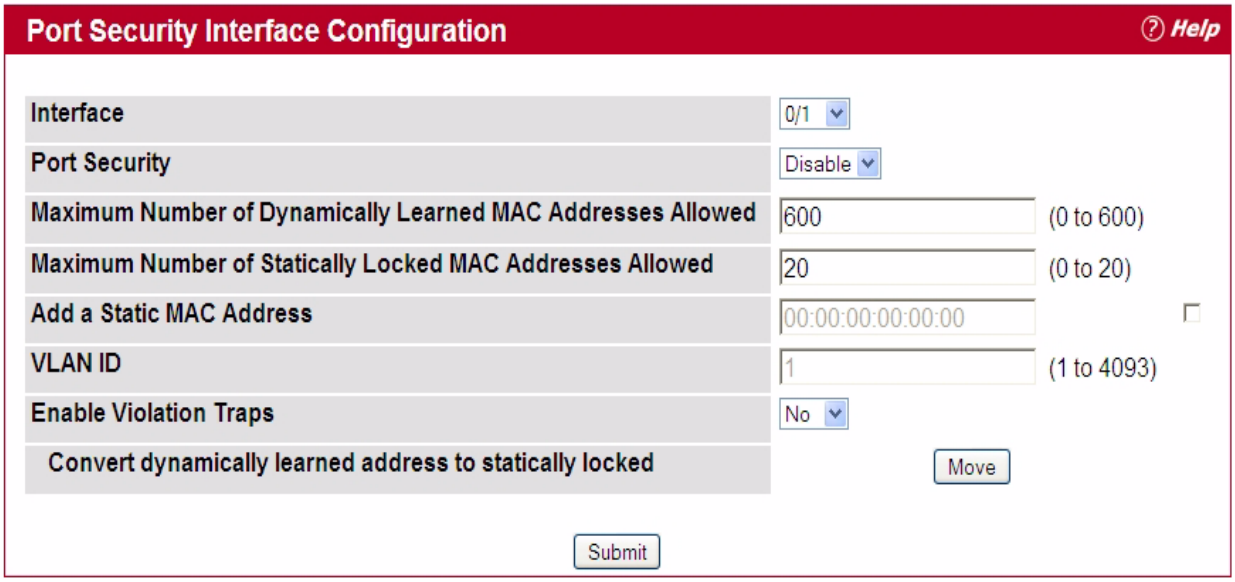

Figure 3-82: Port Security Interface Configuration

Table 3-78: Port Security Interface Configuration Fields

Field	Description
Slot/Port	Select the physical interface or the LAG on which to configure port security information.
Port Security	Determines whether port security is enabled. The default mode is Disable. • **Enable**: Locks the port so that only packets with allowable source MAC addresses can be forwarded. All other packets are discarded. • **Disable**: The port is not locked, so no port security restrictions are applied.
Maximum Number of Dynamically Learned MAC Addresses Allowed	Sets the maximum number of dynamically learned MAC addresses on the selected interface. Once the limit is reached, no more addresses are learned on the port. Any packets with source MAC addresses that were not already learned are discarded. You can effectively disable dynamic locking by setting the number of allowable dynamic entries to zero.
Maximum Number of Statically Locked MAC Addresses Allowed	Sets the maximum number of statically locked MAC addresses on the selected interface.
Add a Static MAC Address	Adds a MAC address to the list of statically locked MAC addresses for the selected interface. Only packets with an allowable source MAC address can be forwarded.
VLAN ID	Adds a corresponding VLAN ID for the MAC Address being added to the list of statically locked MAC addresses for the selected interface.
Enable Violation Traps	Enables or disables the sending of new violation traps designating when a packet with a disallowed MAC address is received on a locked port. Value is No by default.
Convert dynamically learned address to static locked	When you click Move, all the dynamically learned entries on this interface are added to the static MAC address list for this interface. After moving them, you can view them in the Port Security Static page.

If you make any changes to the page, click **Submit** to apply the new settings to the system.

3.20.3 Port Security Static

Use the Port Security Static page to view static MAC addresses configured on an interface.

To access the Port Security Static page, click **Switching > Port Security > Port Security Static** in the navigation tree.

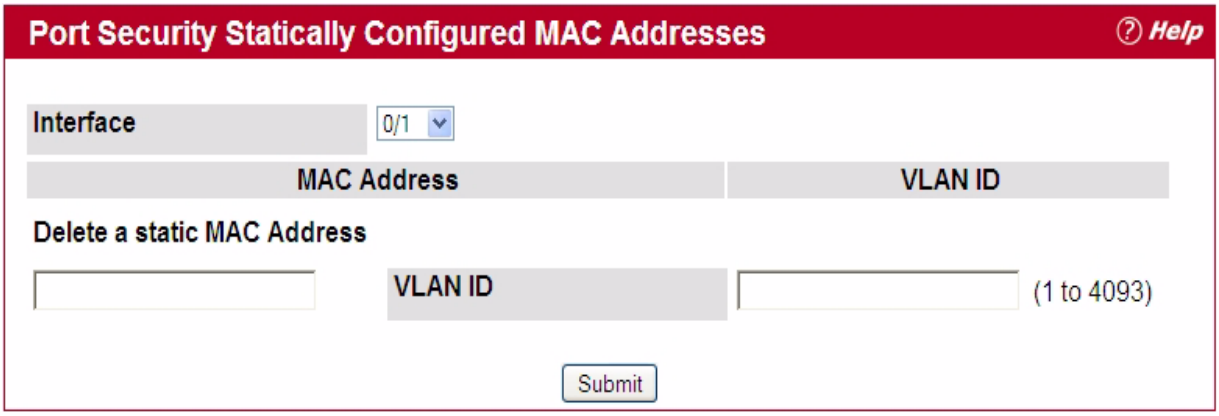

Figure 3-83: Port Security Static

Table 3-79: Port Security Static Fields

Field	Description
Slot/Port	Select the physical interface or the LAG on which to view the dynamically learned MAC addresses.
MAC Address	This column lists the static MAC addresses, if any, configured on the selected port.
VLAN ID	Displays the VLAN ID corresponding to the statically configured MAC address.
Delete a static MAC Address	Enter the address of the statically configured MAC address to delete. All MAC addresses that are available to be deleted appear in the MAC Address – VLAN ID table.
VLAN ID	Enter the VLAN ID that corresponds to the statically configured MAC address to delete.

After you enter the MAC address and VLAN ID of the statically configured MAC address to delete, click **Submit** to remove the MAC address from the port and apply the new settings to the system. The screen refreshes, and the MAC address no longer appears in the table on the page.

3.20.4 Port Security Dynamic

Use the Port Security Dynamic page to view a table with the dynamically learned MAC addresses on an interface. With dynamic locking, MAC addresses are learned on a “first arrival” basis. You specify how many addresses can be learned on the locked port.

To access the Port Security Dynamic page, click **Switching > Port Security > Port Security Dynamic** in the navigation tree.

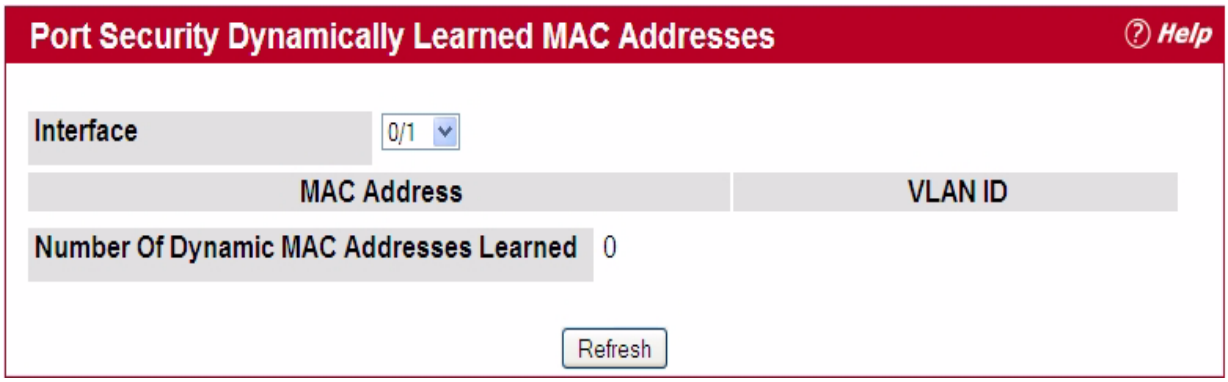

Figure 3-84: Port Security Dynamic

Table 3-80: Port Security Dynamic Fields

Field	Description
Slot/Port	Select the physical interface or the LAG on which to view the dynamically learned MAC addresses.
MAC Address	This column lists the dynamically learned MAC addresses, if any, on the selected port.
VLAN ID	Displays the VLAN ID corresponding to the dynamically learned MAC address.

3.20.5 Port Security Violation Status

Use the Port Security Violation Status page to enable or disable the port security feature on your switch.

To access the Port Security Violation Status page, click **Switching > Port Security > Port Security Violation Status** in the navigation tree.

Figure 3-85: Port Security Violation Status

Table 3-81: Port Security Violation Status Fields

Field	Description
Slot/Port	Select the physical interface or the LAG on which to view security violation information.
Last Violation MAC Address	Displays the source MAC address of the last packet that was discarded at a locked port.
VLAN ID	Displays the VLAN ID corresponding to the Last Violation MAC address.

3.21 Managing LLDP

The IEEE 802.1AB defined standard, Link Layer Discovery Protocol (LLDP), allows stations residing on an 802 LAN to advertise major capabilities and physical descriptions. This information is viewed by a network manager to identify system topology and detect bad configurations on the LAN.

LLDP is a one-way protocol; there are no request/response sequences. Information is advertised by stations implementing the transmit function, and is received and processed by stations implementing the receive function. The transmit and receive functions can be enabled/disabled separately per port. By default, both transmit and receive are disabled on all ports. The application is responsible for starting each transmit and receive state machine appropriately, based on the configured status and operational state of the port.

FASTPATH allows LLDP to have multiple LLDP neighbors per interface. The number of such neighbors is limited by the memory constraints. A product-specific constant defines the maximum number of neighbors supported by the switch. There is no restriction on the number of neighbors supported on a per LLDP port. If all the remote entries on the switch are filled up, the new neighbors are ignored. In case of multiple VOIP devices on a single interface, the 802.1ab component sends the Voice VLAN configuration to all the VoIP devices.

The LLDP folder contains links to the following page:

- Global Configuration
- Interface Configuration
- Interface Summary
- Statistics
- Local Device Information
- Local Device Summary
- Remote Device Information
- Remote Device Summary
- LLDP-MED

3.21.1 Global Configuration

Use the LLDP Global Configuration page to specify LLDP parameters that are applied to the switch.

To display the LLDP Global Configuration page, click **Switching > LLDP > Global Configuration** in the navigation tree.

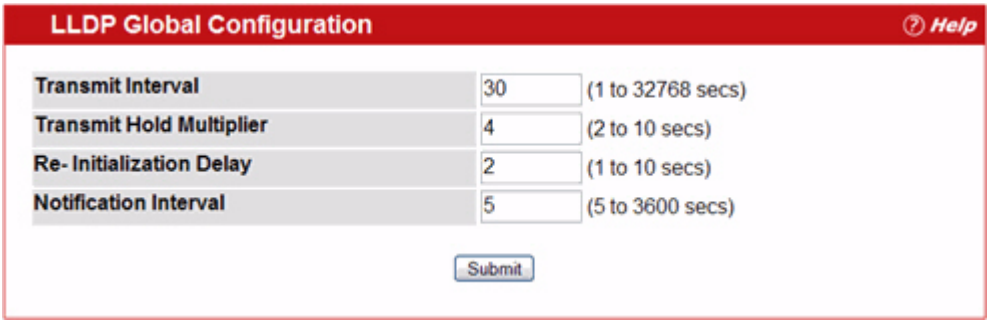

Figure 3-86: LLDP Global Configuration

Table 3-82: LLDP Global Configuration Fields

Field	Description
Transmit Interval	Specifies the interval at which LLDP frames are transmitted. The default is 30 seconds, and the valid range is 1-32768 seconds.
Transmit Hold Multiplier	Specifies multiplier on the transmit interval to assign to TTL. The default is 4, and the range is 2-10.
Re-Initialization Delay	Specifies the delay before a re-initialization. The default is 2 seconds, and the range is 1-10 seconds.
Notification Interval	Limits the transmission of notifications. The default is 5 seconds, and the range is 5-3600 seconds.

If you make any changes to the page, click **Submit** to apply the new settings to the system.

3.21.2 Interface Configuration

Use the LLDP Interface Configuration page to specify LLDP parameters that are applied to a specific interface.

To display the LLDP Interface Configuration page, click **Switching** > **LLDP** > **Interface Configuration** in the navigation tree.

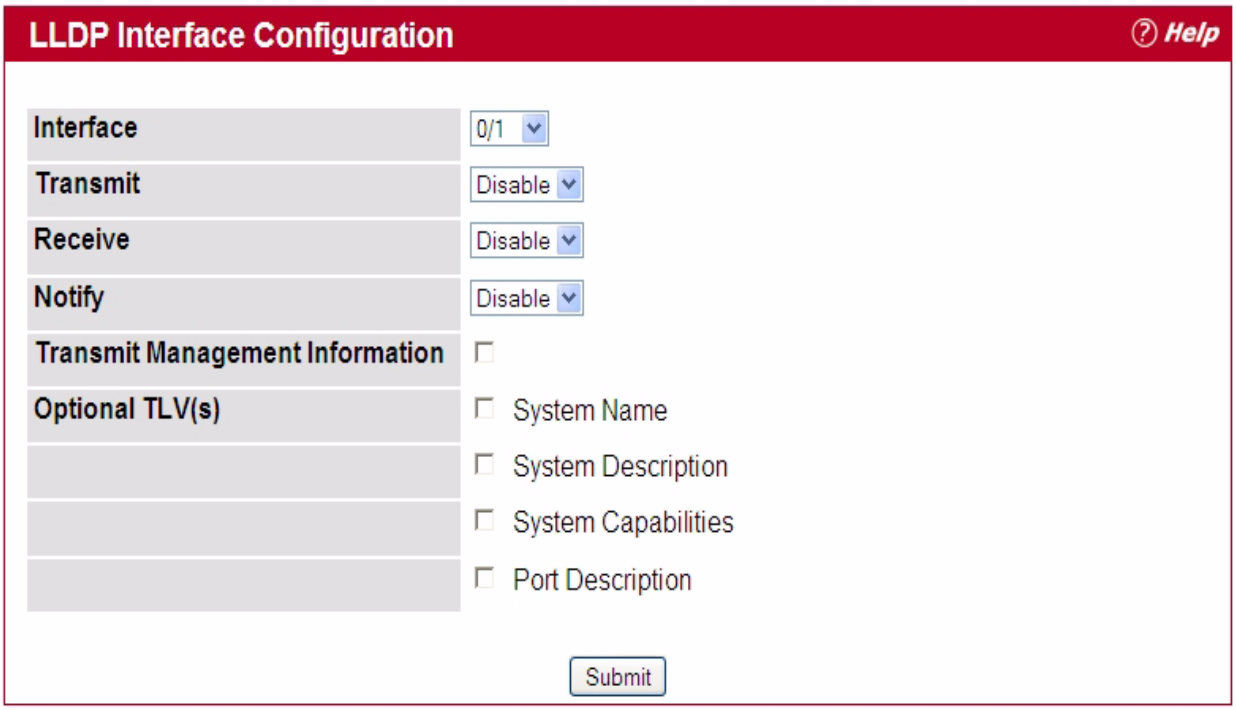

Figure 3-87: LLDP Interface Configuration

Table 3-83: LLDP Interface Configuration Fields

Field	Description
Interface	Specifies the port to be affected by these parameters.
Transmit	Enables or disables the transmission of LLDP protocol data units (PDUs). The default is disabled.
Receive	Enables or disables the ability of the port to receive LLDP PDUs. The default is disabled.
Notify	When notifications are enabled, LLDP interacts with the Trap Manager to notify subscribers of remote data change statistics. The default is disabled.
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.
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. • 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.

If you make any changes to the page, click **Submit** to apply the new settings to the system.

3.21.3 Interface Summary

Use the LLDP Interface Summary page to view the LLDP parameters configured on each physical port on the system.

To display the LLDP Interface Summary page, click **Switching** > **LLDP** > **Interface Summary** in the navigation tree.

LLDP Interface Summary Help

Interface	Link Status	Transmit	Receive	Notify	Optional TLV(s)	Transmit Management Information
0/1	Down	Disable	Disable	Disable		No
0/2	Down	Disable	Disable	Disable		No
0/3	Down	Disable	Disable	Disable		No
0/4	Down	Disable	Disable	Disable		No
0/5	Down	Disable	Disable	Disable		No
0/6	Down	Disable	Disable	Disable		No
0/7	Down	Disable	Disable	Disable		No
0/8	Down	Disable	Disable	Disable		No
0/9	Down	Disable	Disable	Disable		No
0/10	Down	Disable	Disable	Disable		No
0/11	Down	Disable	Disable	Disable		No
0/12	Down	Disable	Disable	Disable		No
0/13	Down	Disable	Disable	Disable		No
0/14	Down	Disable	Disable	Disable		No
0/15	Down	Disable	Disable	Disable		No
0/16	Down	Disable	Disable	Disable		No

Refresh

Figure 3-88: LLDP Interface Summary

Table 3-84: LLDP Interface Summary Fields

Field	Description
Interface	Displays all the ports on which LLDP-802.1AB can be configured.
Link Status	Displays whether the link status of the ports is up or down.
Transmit	Displays the LLDP-802.1AB transmit mode of the interface.

Table 3-84: LLDP Interface Summary Fields (Continued)

Field	Description
Receive	Displays the LLDP-802.1AB receive mode of the interface.
Notify	Displays the LLDP-802.1AB notification mode of the interface.
Optional TLV(s)	Shows the LLDP-802.1AB optional type-length values (TLV) that are included. If no TVLs are sent, the entry is blank. The field can contain one or more of the following TVLs. • System Name • System Capabilities • System Description • Port Description.
Transmit Management Information	Shows whether the management address is transmitted in the LLDP frames.

To update the page with the latest data, click **Refresh**.

3.21.4 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 tree.

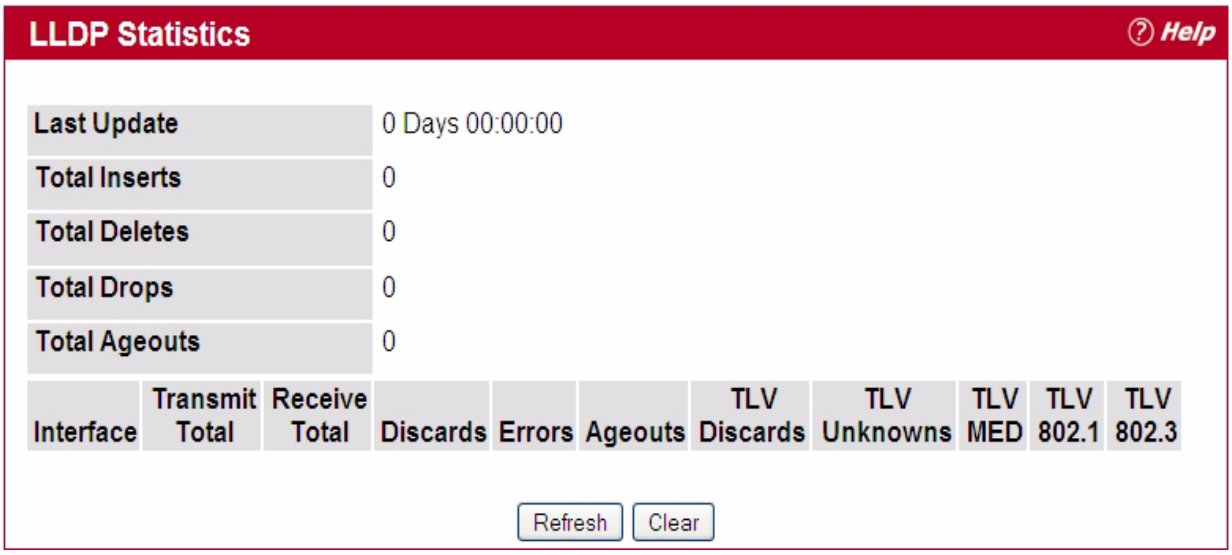

Figure 3-89: LLDP Statistics

Table 3-85: LLDP Statistics Fields

Field	Description
	System-wide Statistics
Last Update	Displays the time when an entry was created, modified, or deleted in the tables associated with the remote systems.
Total Inserts	Displays the number of times a complete set of information advertised by a particular MAC Service Access Point (MSAP) has been inserted into the tables associated with the remote systems.
Total Deletes	Displays the number of times a complete set of information advertised by a particular MAC Service Access Point (MSAP) has been deleted from the tables associated with the remote systems.
Total Drops	Displays the number of times a complete set of information advertised by a particular MAC Service Access Point (MSAP) could not be entered into tables associated with the remote systems because of insufficient resources.
Total Ageouts	Displays the number of times a complete set of information advertised by a particular MAC Service Access Point (MSAP) has been deleted from tables associated with the remote systems because the information timelines interval has expired.
	Port Statistics
Interface	Displays the slot/port for 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.

- Click **Refresh** to update the page with the most current information.
- Click **Clear** to clear the LLDP statistics of all the interfaces.

3.21.5 Local Device Information

Use the LLDP Local Device Information page to view the data that each port advertises through LLDP.

To display the LLDP Local Device Information page, click **Switching** > **LLDP** > **Local Device Information** in the navigation tree.

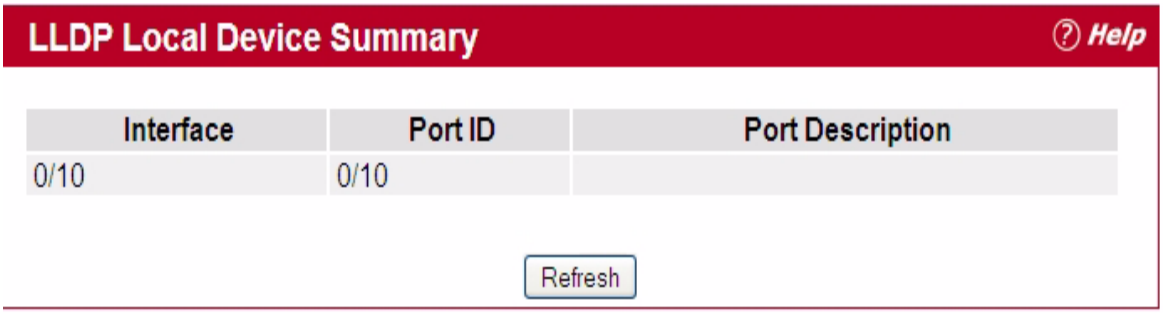

Figure 3-90: LLDP Local Device Information

Table 3-86: LLDP Local Device Information Fields

Field	Description
Interface	Select from the list of all the ports on which LLDP-802.1AB frames can be transmitted.
Chassis ID Subtype	Displays the string that describes the source of the chassis identifier.
Chassis ID	Displays the string value used to identify the chassis component associated with the local system.
Port ID Subtype	Displays the string describing the source of the port identifier.
Port ID	Identifies the physical address of the port.
System Name	Displays the system name of the local system.
System Description	Displays the description of the selected port associated with the local system.
Port Description	Displays the user-defined description of the port.
System Capabilities Supported	Displays the system capabilities of the local system.
System Capabilities Enabled	Displays the system capabilities of the local system which are supported and enabled.
Management Address	Displays the advertised management address of the local system.
Management Address Type	Specifies the type of the management address.

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

3.21.6 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 LLDP Local Device Summary page, click **Switching** > **LLDP** > **Local Device Summary** in the navigation tree.

Figure 3-91: LLDP Local Device Summary

Table 3-87: LLDP Local Device Summary Columns

Field	Description
Interface	Displays the slot/port on which LLDP-802.1AB frames can be transmitted.
Port ID	Displays the string describing the source of the port identifier.
Port Description	Displays the description of the port associated with the local system.

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

3.21.7 Remote Device Information

Use the LLDP Remote Device Information page to view the data that a specified interface has received from other LLDP-enabled systems.

To display the LLDP Remote Device Information page, click **Switching > LLDP > Remote Device Information** in the navigation tree.

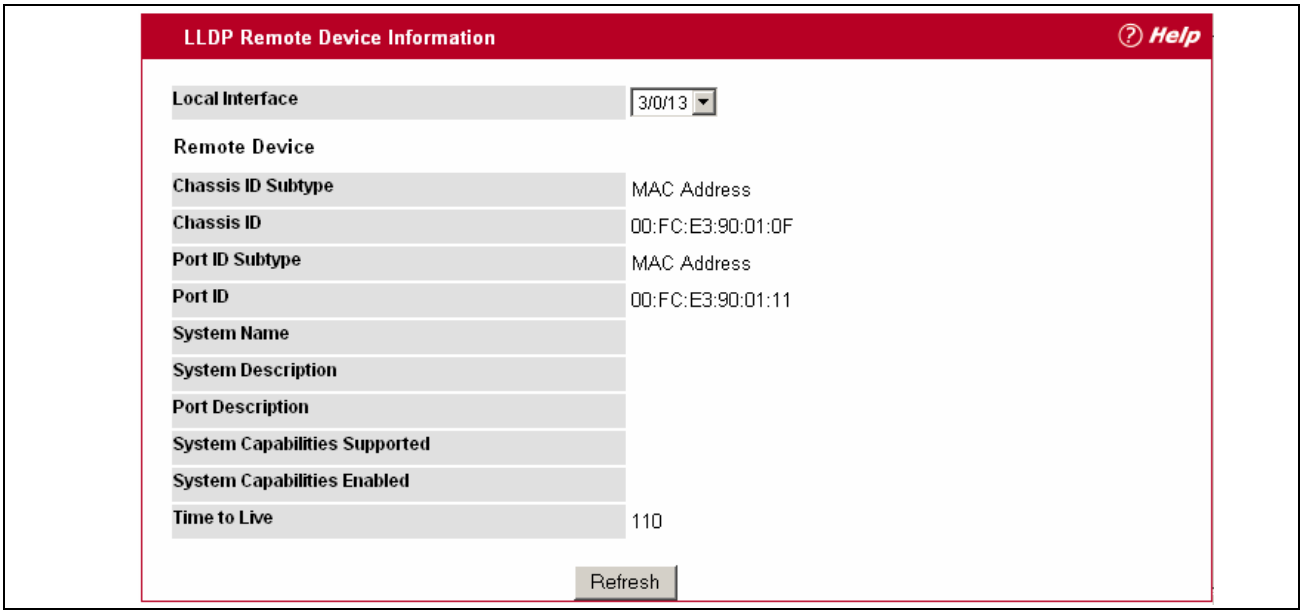

Figure 3-92: LLDP Remote Device Information

Table 3-88: LLDP Remote Device Information Fields

Field	Description
Local Interface	Select the slot/port on the local system to display the LLDP information it has received. **Note:** If no LLDP data has been received on the select interface, a message stating so displays. If the selected interface has received LLDP information from a remote device, the following fields display:
Remote ID	Displays the remote client identifier assigned to the remote system.
Chassis ID Subtype	Identifies the type of data displayed in the **Chassis ID** field on the remote system.
Chassis ID	Identifies the chassis component associated with the remote system.
Port ID Subtype	Identifies the type of data displayed in the remote system's **Port ID** field.
Port ID	Identifies the physical address of the port on the remote system from which the data was sent.
System Name	Identifies the system name of the remote system.
System Description	Displays the description of the selected port associated with the remote system.
Port Description	Displays the user-defined description of the port.
System Capabilities Supported	Displays the system capabilities of the remote system.
System Capabilities Enabled	Displays the system capabilities of the remote system which are supported and enabled.
Time to Live	Displays the Time to Live value in seconds of the received remote entry.
Management Address	Displays the advertised management address of the remote system.
Management Address Type	Displays the type of the management address.

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

3.21.8 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 Device Summary** in the navigation tree.

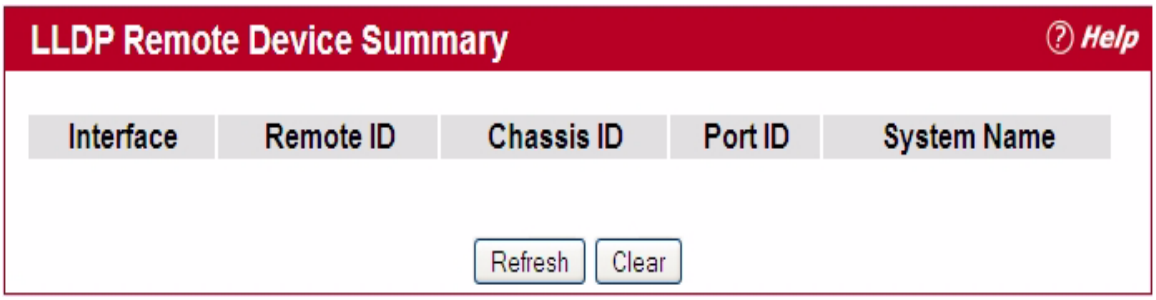

Figure 3-93: LLDP Remote Device Summary

Table 3-89: LLDP Remote Device Summary Columns

Field	Description
Local Interface	Shows the slot/port on the local system that can receive LLDP frames advertised by a remote system.
Chassis ID	Identifies the chassis component associated with the remote system.
Port ID	Identifies the physical address of the port on the remote device that sent the LLDP data.
Remote ID	Shows the remote client identifier assigned to the remote system.
System Name	Shows the system name of the remote device. If the system name is not configured, the field is blank.

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

3.21.9 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).

The LLDP-MED folder provides access to the following pages:

- LLDP-MED Global Configuration
- LLDP-MED Interface Configuration
- LLDP-MED Interface Summary
- LLDP Local Device Information
- LLDP-MED Remote Device Information

3.21.9.1 LLDP-MED Global Configuration

Use this page to set global parameters for LLDP-MED operation. To display this page, click **Switching > LLDP > LLDP-MED > Global Configuration** in the navigation tree.

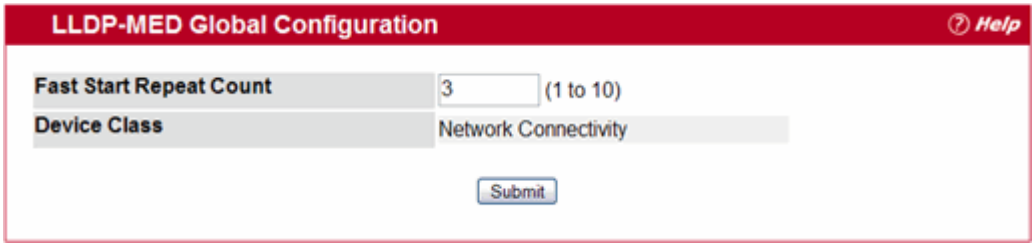

Figure 3-94: LLDP Global Configuration

Table 3-90: LLDP Global Configuration Fields

Field	Description
Fast Start Repeat Count	Specifies the number of LLDP PDUs that will be transmitted when the protocol is enabled. The range is from (1 to 10). The default value s 3.
Device Class	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/router, IEEE 802.1 bridge, IEEE 802.11 wireless access point, etc.

Click **Submit** to updated the switch. The changes take effect but will not be retained across a power cycle unless a save is performed.

3.21.9.2 LLDP-MED Interface Configuration

Use this page to enable LLDP-MED mode on an interface and configure its properties. To display this page, click **Switching > LLDP > LLDP-MED > Interface Configuration** in the navigation tree.

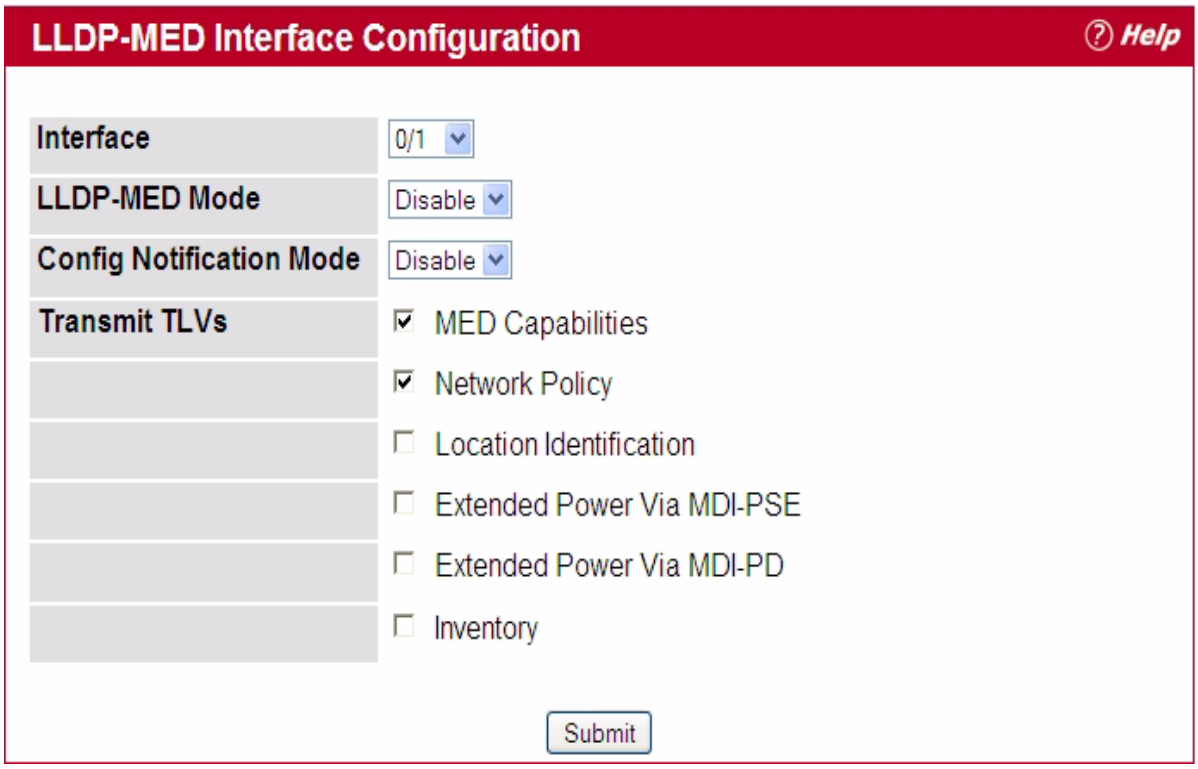

Figure 3-95: LLDP-MED Interface Configure

Table 3-91: LLDP-MED Interface Configuration Fields

Field	Description
Interface	Selects the port that you want to configure LLDP-MED - 802.1AB on. You can select **All** to configure all interfaces on the DUT with the same properties. To view the summary of all interfaces, refer to the 3.21.9.3 LLDP-MED Interface Summary. The Interface Configuration page will not be able to display the summary of 'All' interfaces. The summary of individual interfaces is visible from the Interface Configuration page. The Interface Configuration page for the 'All' option will always display the LLDP-MED mode and notification mode as 'disabled' and checkboxes for 'Transmit TLVs' will always be unchecked.
LLDP-MED Mode	Enables or disables LLDP-MED mode for the selected interface. By enabling MED, you will be effectively enabling the transmit and receive function of LLDP.
Config Notification Mode	Enables or disables LLDP-MED topology change notification mode for the selected interface.
Transmit TLVs	Specifies which optional type length values (TLVs) in the LLDP-MED will be transmitted in the LLDP PDUs frames for the selected interface: • MED Capabilities: Transmits the capabilities TLV in LLDP frames. • Network Policy: Transmits the network policy TLV in LLDP frames. • Location Identification: Transmits the location TLV in LLDP frames. • Extended Power via MDI - PSE: Transmits the extended PSE TLV in LLDP frames. • Extended Power via MDI - PD: Transmits the extended PD TLV in LLDP frames. • Inventory: Transmits the inventory TLV in LLDP frames.

Click **Submit** to send the updated configuration to the switch. These changes take effect immediately but will not be retained across a power cycle unless a save is performed.

3.21.9.3 LLDP-MED Interface Summary

This page lists each switch interface and its LLDP configuration status. To display this page, click **Switching > LLDP > LLDP-MED > Interface Summary** in the navigation tree.

LLDP-MED Interface Summary Help

Interface	Link Status	MED Status	Operational Status	Notification Status	Transmit TLVs
0/1	Down	Disable	Disable	Disable	Capabilities Network Policy
0/2	Down	Disable	Disable	Disable	Capabilities Network Policy
0/3	Down	Disable	Disable	Disable	Capabilities Network Policy
0/4	Down	Disable	Disable	Disable	Capabilities Network Policy
0/5	Down	Disable	Disable	Disable	Capabilities Network Policy
0/6	Down	Disable	Disable	Disable	Capabilities Network Policy
0/7	Down	Disable	Disable	Disable	Capabilities Network Policy
0/8	Down	Disable	Disable	Disable	Capabilities Network Policy
0/9	Down	Disable	Disable	Disable	Capabilities Network Policy
0/10	Down	Disable	Disable	Disable	Capabilities Network Policy
0/11	Down	Disable	Disable	Disable	Capabilities Network Policy
0/12	Down	Disable	Disable	Disable	Capabilities Network Policy
0/13	Down	Disable	Disable	Disable	Capabilities Network Policy
0/14	Down	Disable	Disable	Disable	Capabilities Network Policy
0/15	Down	Disable	Disable	Disable	Capabilities Network Policy
0/16	Down	Disable	Disable	Disable	Capabilities Network Policy

Refresh

Figure 3-96: LLDP-MED Interface Summary

Table 3-92: LLDP-MED Interface Summary Fields

Field	Description
Interface	Specifies all the ports on which LLDP-MED can be configured.
Link Status	Specifies the link status of the ports as Up/Down.
MED Status	Specifies the transmit and/or receive LLDP-MED mode is enabled or disabled on this interface.
Operational Status	Specifies whether the interface will transmit TLVs.
Notification Status	Specifies the LLDP-MED topology notification mode of the interface.
Transmit TLVs	Specifies the LLDP-MED transmit TLV(s) that are included.

Click **Refresh** to update the page with the latest information from the router.

3.21.9.4 LLDP Local Device Information

This page displays information on LLPD-MED information advertised on the selected local interface. To display this page, click **Switching > LLDP > LLDP-MED > Local Device Information** in the navigation tree.

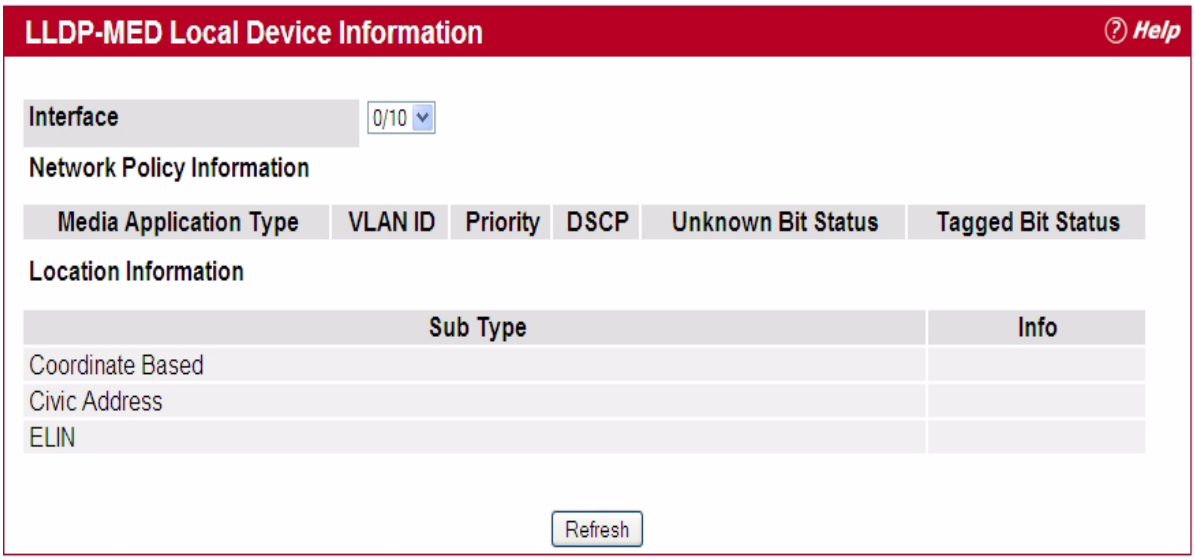

Figure 3-97: LLPD-MED Local Device Information

Table 3-93: LLPD-MED Local Device Information Fields

Field	Description
Interface	Select from the list of all the ports on which LLDP-MED frames can be transmitted.
Network Policy Information	Specifies if network policy TLV is present in the LLDP frames: • **Media Application Type**: Specifies the application type. Types of application types are unknown, voicesignaling, guestvoice, guestvoicesignalling, softphonevoice, videoconferencing, streammingvideo, vidoesignalling. Each application type that is received has the VLAN ID, priority, DSCP, tagged bit status and unknown bit status. A port may receive one or many such application types. This information is only displayed when a network policy TLV has been transmitted. • **Vlan Id**: Specifies the VLAN id associated with a particular policy type. • **Priority**: Specifies the priority associated with a particular policy type. • **DSCP**: Specifies the DSCP associated with a particular policy type. • **Unknown Bit Status**: Specifies the unknown bit associated with a particular policy type. • **Tagged Bit Status**: Specifies the tagged bit associated with a particular policy type.
Inventory	Specifies the inventory TLV present in LLDP frames: • **Hardware Revisions.** Specifies hardware version. • **Firmware Revisions.** Specifies firmware version. • **Software Revisions.** Specifies software version. • **Serial Number.** Specifies serial number. • **Manufacturer Name**. Specifies manufacturer's name. • **Model Name.** Specifies model name. • **Asset ID.** Specifies asset ID.
Location Information	Specifies if location TLV is present in LLDP frames: • **Sub Type**: Specifies type of location information. • **Location Information**: Specifies the location information as a string for given type of location ID.
Extended PoE	Specifies if local device is a PoE device. • **Device Type**. Specifies power device type.
Extended PoE PSE	Specifies if extended PSE TLV is present in LLDP frame: • **Available**: Specifies available power sourcing equipment's power value in tenths of watts on the port of local device. • **Source**: Specifies power source of this port. • **Priority**: Specifies PSE port power priority.
Extended PoE PD	Specifies if extended PD TLV is present in LLDP frame. • **Required**: Specifies required power device power value in tenths of watts on the port of local device. • **Source**: Specifies power source of this port. • **Priority**: Specifies PD port power priority.

Click **Refresh** to update the page with the latest information from the router.

3.21.9.5 LLDP-MED Remote Device Information

This page displays information on LLPD-MED information received from remote clients on the selected local interface. To display this page, click **Switching > LLDP > LLDP-MED > Remote Device Information** in the navigation tree.

Figure 3-98: LLDP Remote Device Information

Table 3-94: LLPD-MED Local Device Information Fields

Field	Description
Local Interface	Specifies the list of all the ports on which LLDP-MED is enabled.
Remote ID	Specifies the remote client identifier assigned to the remote system.
Capability Information	Specifies the supported and enabled capabilities that were received in MED TLV on this port: • **Supported Capabilities**: Specifies supported capabilities that were received in MED TLV on this port. • **Enabled Capabilities**: Specifies enabled capabilities that were received in MED TLV on this port. • **Device Class**: Specifies device class as advertised by the device remotely connected to the port.
Network Policy Information	Specifies if network policy TLV is received in the LLDP frames on this port: • **Media Application Type**: Specifies the application type. Types of application types are unknown, voicesignaling, guestvoice, guestvoicesignalling, softphonevoice, videoconferencing, streammingvideo, vidoesignalling. Each application type that is received has the VLAN ID, priority, DSCP, tagged bit status and unknown bit status. A port may receive one or many such application types. This information is displayed only when a network policy TLV has been received on this port. • **Vlan ID**: Specifies the VLAN ID associated with a particular policy type. • **Priority**: Specifies the priority associated with a particular policy type. • **DSCP**: Specifies the DSCP associated with a particular policy type. • **Unknown Bit Status**: Specifies the unknown bit associated with a particular policy type. • **Tagged Bit Status**: Specifies the tagged bit associated with a particular policy type.
Inventory	Specifies the inventory TLV is received in LLDP frames on this port: • **Hardware Revisions.** Specifies hardware version of the remote device. • **Firmware Revisions.** Specifies firmware version of the remote device. • **Software Revisions.** Specifies software version of the remote device. • **Serial Number.** Specifies serial number of the remote device. • **Manufacturer Name.** Specifies manufacturer's name of the remote device. • **Model Name.** Specifies model name of the remote device. • **Asset ID.** Specifies asset ID of the remote device.

Table 3-94: LLPD-MED Local Device Information Fields (Continued)

Field	Description
Location Information	Specifies if location TLV is received in LLDP frames on this port. • **Sub Type**: Specifies type of location information. • **Location Information**: Specifies the location information as a string for given type of location ID.
Extended PoE	Specifies if remote device is a PoE device. • **Device Type**. Specifies the remote device's PoE device type connected to this port.
Extended PoE PSE	Specifies if extended PSE TLV is received in LLDP frame on this port: • **Available**: Specifies the remote port's power sourcing equipment's (PSE) power value in tenths of watts. • **Source**: Specifies the remote port's PSE power source. • **Priority**: Specifies the remote port's PSE power priority.
Extended PoE PD	Specifies if extended PD TLV is received in LLDP frame on this port. • **Required**: Specifies the remote port's power device power requirement. • **Source**: Specifies the remote port's PD power source. • **Priority**: Specifies the remote port's PD power priority.

Click **Refresh** to update the page with the latest information from the router.

4 Configuring Routing

FASTPATH supports IP routing. Use the links in the Routing navigation tree folder to manage routing on the system. This section contains the following information:

- Configuring ARP
- Configuring Global IP Settings
- Configuring OSPF
- Managing the BOOTP/DHCP Relay Agent
- IP Helper
- Configuring RIP
- Router Discovery
- Router
- VLAN Routing
- Virtual Router Redundancy Protocol (VRRP)
- Tunnels
- Loopback Interfaces

When a packet enters the switch, the destination MAC address is checked to see if it matches any of the configured routing interfaces. If it does, then the silicon searches the host table for a matching destination IP address. If an entry is found, then the packet is routed to the host. If there is not a matching entry, then the switch performs a longest prefix match on the destination IP address. If an entry is found, then the packet is routed to the next hop. If there is no match, then the packet is routed to the next hop specified in the default route. If there is no default route configured, then the packet is passed to the 6200 series software to be handled appropriately.

The routing table can have entries added either statically by the administrator or dynamically via a routing protocol. The host table can have entries added either statically by the administrator or dynamically via ARP.

4.1 Configuring ARP

The ARP protocol associates a layer 2 MAC address with a layer 3 IPv4 address. FASTPATH 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 folder contains links to the following web pages that configure and display ARP detail:

- ARP Create
- ARP Table Configuration

4.1.1 ARP Create

Use the ARP Create page to add an entry to the Address Resolution Protocol table.

To display the page, click **Routing > ARP > ARP Create** in the navigation tree.

Figure 4-1: ARP Create

Table 4-1: ARP Create Fields

Field	Description
IP Address	Enter the IP address you want to add. It must be the IP address of a device on a subnet attached to one of the switch's existing routing interfaces.
MAC Address	The unicast MAC address of the device. Enter the address as six two-digit hexadecimal numbers separated by colons, for example 00:06:29:32:81:40.

After you enter an IP address and the associated MAC address, click **Submit** to apply the changes to the system and create the entry in the ARP table.

4.1.2 ARP Table Configuration

Use this page to change the configuration parameters for the Address Resolution Protocol Table. You can also use this screen to display the contents of the table.

To display the page, click **Routing > ARP > ARP Create** in the navigation tree.

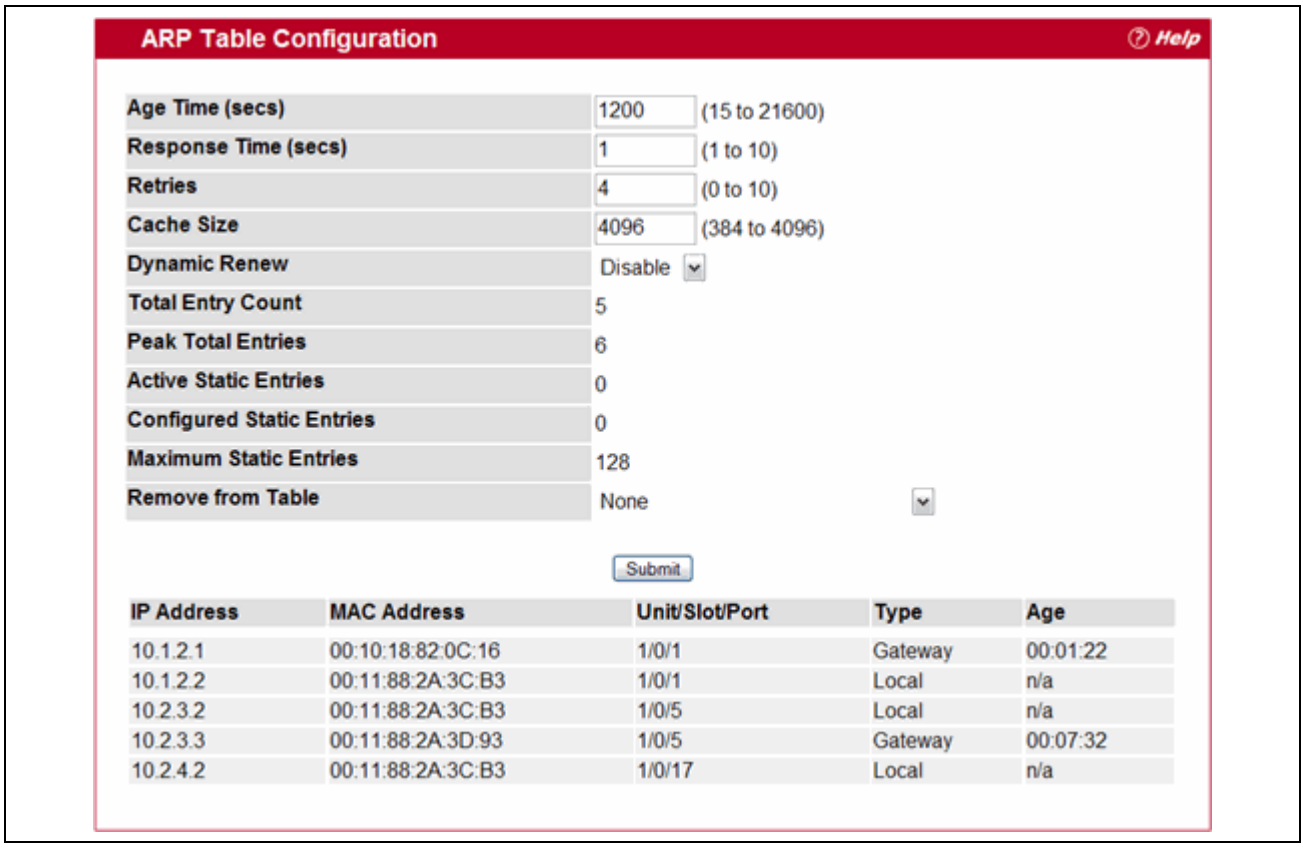

Figure 4-2: ARP Table Configuration

Table 4-2: ARP Table Configuration Fields

Field	Description
Age Time (secs)	Enter the value you want the switch to use for the ARP entry ageout time. You must enter a valid integer, which represents the number of seconds it takes for an ARP entry to age out. The range for this field is 15 to 21600 seconds. The default value for Age Time is 1200 seconds.
Response Time (secs)	Enter the value you want the switch to use for the ARP response timeout. You must enter a valid integer, which represents the number of seconds the switch waits for a response to an ARP request. The range for this field is 1 to 10 seconds. The default value for Response Time is 1 second.
Retries	Enter an integer which specifies the maximum number of times an ARP request is retried. The range for this field is 0 to 10. The default value for Retries is 4.
Cache Size	Enter an integer which specifies the maximum number of entries for the ARP cache. The range for this field is platform-dependent. The default value for Cache Size is 896.
Dynamic Renew	This controls whether the ARP component automatically attempts to renew ARP Entries of type Dynamic when they age out. The default setting is Disable.
Total Entry Count	Total number of entries in the ARP table.
Peak Total Entries	Highest value reached by Total Entry Count. This counter value is restarted whenever the ARP table Cache Size value is changed.

Table 4-2: ARP Table Configuration Fields (Continued)

Field	Description
Active Static Entries	Total number of active static entries in the ARP table.
Configured Static Entries	Total number of configured static entries in the ARP table.
Maximum Static Entries	Maximum number of static entries that can be defined.
Remove from Table	Allows you to remove certain entries from the ARP Table. The choices listed specify the type of ARP Entry to be deleted: • All Dynamic Entries • All Dynamic and Gateway Entries • Specific Dynamic Gateway Entry • Specific Static Entry

The ARP Table displays at the bottom of the page, and contains the following fields:

Table 4-3: ARP Table Fields

Field	Description
IP Address	The IP address of a device on a subnet attached to one of the switch's routing interfaces.
MAC Address	The unicast MAC address for the device. The format is six two-digit hexadecimal numbers separated by colons, for example 00:06:29:32:81:40.
Slot/Port	The routing interface associated with the ARP entry.
Type	The type of the ARP entry.
Age	Age since the entry was last refreshed in the ARP Table. The format is hh:mm:ss

If you make any changes to the page, click **Submit** to apply the changes to the system.

4.2 Configuring Global IP Settings

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

- IP Configuration
- IP Statistics
- IP Interface Configuration

4.2.1 IP Configuration

Use the IP Configuration page to configure routing parameters for the switch as opposed to an interface.

To display the page, click **Routing > IP > Configuration** in the navigation tree.

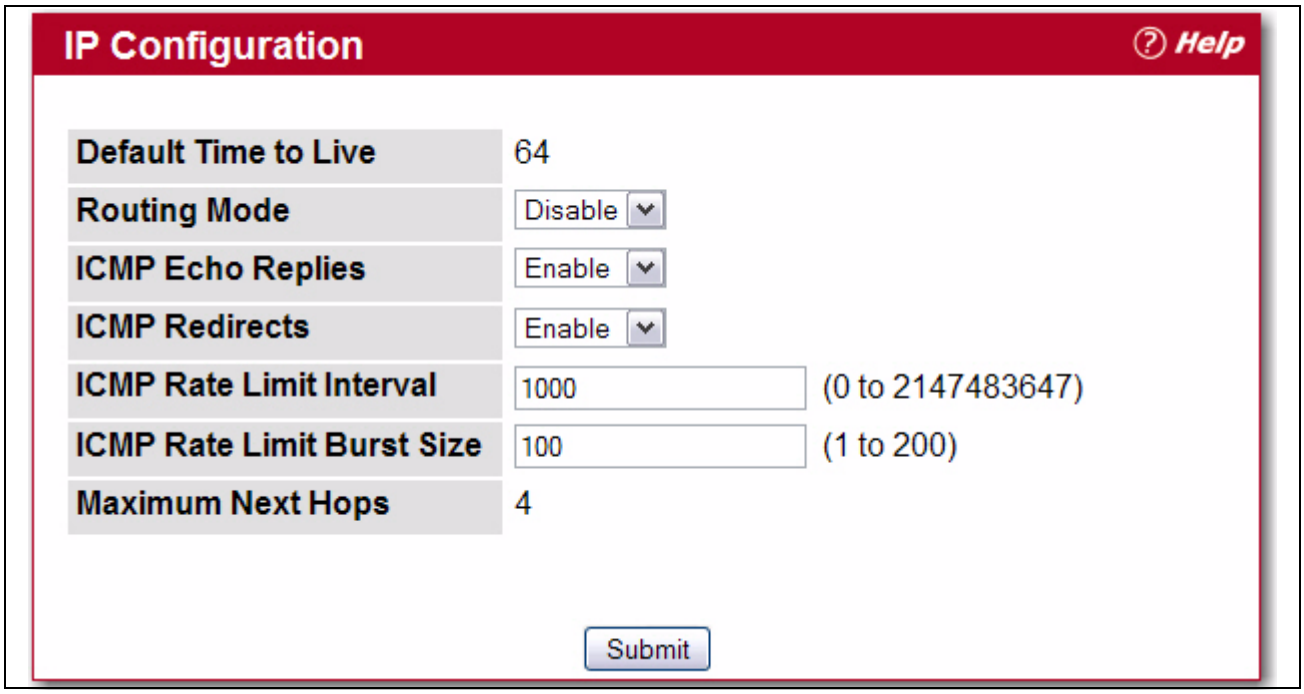

Figure 4-3: IP Configuration

Table 4-4: IP Configuration Fields

Field	Description
Default Time to Live	The default value inserted into the Time-To-Live field of the IP header of datagrams originated by the switch, if a TTL value is not supplied by the transport layer protocol.
Routing Mode	Select **Enable** or **Disable** from the dropdown menu. You must enable routing for the switch before you can route through any of the interfaces. Routing is also enabled or disabled per VLAN interface. The default value is **Disable**.
ICMP Echo Replies	Select **Enable** or **Disable** from the dropdown 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	If this is enabled globally and on the interface level, then only the router can send ICMP redirects.
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.
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.
Maximum Next Hops	The maximum number of hops supported by the switch. This is a read-only value.

If you make any changes to the page, click **Submit** to apply the changes to the system.

4.2.2 IP Statistics

The statistics reported on the IP Statistics page are as specified in RFC 1213.

To display the page, click **Routing > IP > Statistics** in the navigation tree.

> **Note...**
>
> Figure 4-4 does not show all of the fields on the page.

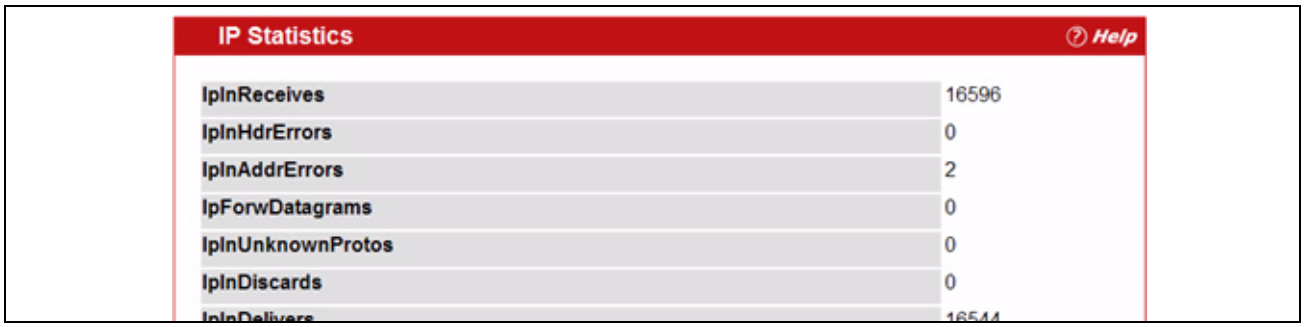

Figure 4-4: IP Statistics

Table 4-5: 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.

Table 4-5: IP Statistics Fields (Continued)

Field	Description
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.
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.

Table 4-5: IP Statistics Fields (Continued)

Field	Description
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.
IcmpOutAddrMaskReps	The number of ICMP Address Mask Reply messages sent.

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

4.2.3 IP Interface Configuration

Use the IP Interface Configuration page to update IP interface data for this switch.

To display the page, click **Routing > IP > Interface Configuration** in the navigation tree.

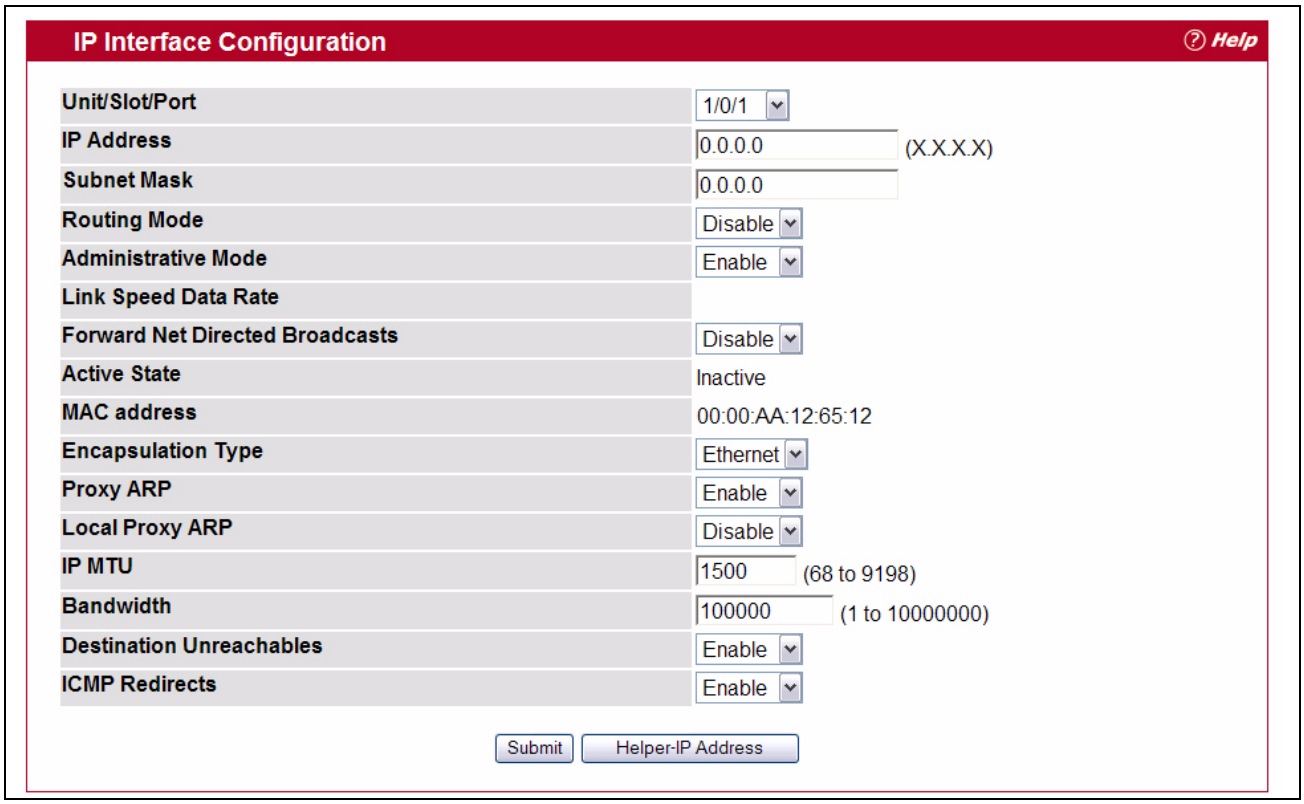

Figure 4-5: IP Interface Configuration

Table 4-6: IP Interface Configuration Fields

Field	Description
Slot/Port	Select the interface to configure from the dropdown menu. The dropdown menu contains logical interfaces, including loopback interfaces and VLAN routing interfaces.
IP Address	Enter the IP address for the interface.
Subnet Mask	Enter the subnet mask for the interface. This is also referred to as the subnet/ network mask, and defines the portion of the interface's IP address that is used to identify the attached network.
Routing Mode	Setting this Enables or Disables routing for an interface. By default, routing is disabled on port-based routing interfaces and enabled on VLAN-based routing interfaces.
Administrative Mode	The Administrative Mode of the interface. The default value is Enable.
Link Speed Data Rate	An integer representing the physical link data rate of the specified interface. This data is valid only for physical interfaces and is measured in Megabits per second (Mbps).
Forward Net Directed Broadcasts	Select how network directed broadcast packets should be handled. If you select Enable from the dropdown menu network directed broadcasts are forwarded. If you select Disable they are dropped. The default value is Disable.
Active State	The state of the specified interface is either Active or Inactive. An interface is considered active if the link is up and it is in forwarding state.
MAC Address	The burned-in physical address of the specified interface. The format is six two-digit hexadecimal numbers separated by colons, for example 00:06:29:32:81:40. This value is valid for physical interfaces. For logical interfaces, such as VLAN routing interfaces, the field displays the system MAC address.
Encapsulation Type	Select the link layer encapsulation type for packets transmitted from the specified interface from the dropdown menu. The possible values are Ethernet and SNAP. The default is Ethernet.
Proxy ARP	Select to Disable or Enable Proxy ARP for the specified interface from the dropdown menu.
Local Proxy ARP	Select to Disable or Enable Local Proxy ARP for the specified interface from the dropdown menu.
IP MTU	The maximum transmission unit (MTU) size of IP packets sent on an interface. Valid range is (68 to 9198). Default value is 1500.
Bandwidth	The configured bandwidth of the interface is specified in Kbps. The OSPF protocol uses this value to compute the link cost of an interface as the ratio of the reference bandwidth to the interface bandwidth. If no bandwidth is configured, the bandwidth defaults to the actual interface bandwidth for port-based routing interfaces and to 10 Mbps for VLAN routing interfaces. This value does not affect the actual speed of an interface.
Destination Unreachables	Specifies the mode of sending ICMP Destination Unreachables on this interface. If this is disabled, then this interface will not send ICMP Destination Unreachables. By default, the Destination Unreachables mode is **Enable**.
ICMP Redirects	The router sends an ICMP Redirect on an interface only if Redirects are enabled both globally and on the interface. By default, the ICMP Redirects mode is **Enable.**

- Click **Submit t**o send the updated configuration to the switch. These changes will not be retained across a power cycle unless a Save configuration is performed.
- Click **Helper-IP Address** to proceed to the Helper Address configuration page.

4.3 Configuring OSPF

The **Routing > OSPF** folder contains the following links to web pages that configure and display OSPF parameters and data:

- OSPF Configuration
- OSPF Area Configuration
- OSPF Stub Area Summary
- OSPF Area Range Configuration
- OSPF Interface Statistics
- OSPF Interface Configuration
- Neighbor Table
- OSPF Neighbor Configuration
- OSPF Link State Database
- OSPF Virtual Link Configuration
- OSPF Virtual Link Summary
- OSPF Route Redistribution Configuration
- OSPF Route Redistribution Summary

4.3.1 OSPF Configuration

Use the OSPF Configuration page to enable OSPF on a router and to configure the related OSPF settings.

To display the page, click **Routing > OSPF > OSPF Configuration** in the navigation menu.

OSPF Configuration Help

Router ID	0.0.0.0	
OSPF Admin Mode	Enable	
ASBR Status	Disabled	
RFC 1583 Compatibility	Enable	
ABR Status		
Opaque LSA Status	Disable	
Exit Overflow Interval (secs)	0	(0 to 2147483647)
SPF DelayTime(secs)	5	(0 to 65535)
SPF HoldTime(secs)	10	(0 to 65535)
External LSA Count		
External LSA Checksum		
AS_OPAQUE LSA Count		
AS_OPAQUE LSA Checksum		
New LSAs Originated		
LSAs Received		
External LSDB Limit	No Limit	(-1(No Limit) to 2147483647)
Default Metric		(1 to 16777214)
Maximum Paths	4	(1 to 4)
AutoCost Reference Bandwidth	100	(1 to 4294967)
Default Passive Setting	Disable	
Default Route Advertise		
Default Information Originate	Disable	
Always	False	
Metric		(0 to 16777214)
Metric Type	External Type 2	

Submit

Figure 4-6: OSPF Configuration

Table 4-7: OSPF Configuration Fields

Field	Description
Router ID	The 32-bit integer in dotted decimal format that uniquely identifies the router within the autonomous system (AS). If you want to change the Router ID you must first disable OSPF. After you set the new Router ID, you must re-enable OSPF to have the change take effect. The default value is 0.0.0.0, although this is not a valid Router ID.
OSPF Admin Mode	Select Enable or Disable from the dropdown menu. If you select Enable OSPF is activated for the switch. The default value is Enable. You must configure a Router ID before OSPF can become operational. You do this on the IP Configuration page or by issuing the CLI following commands: `config router ospf router-id <value>` **Note**: Once OSPF is initialized on the router, it remains initialized until the router is reset.
ASBR Status	A router is an autonomous system boundary router (ASBR) if it is configured to redistribute routes from another protocol.
RFC 1583 Compatibility	Select Enable or Disable from the dropdown menu to specify the preference rules that are used when choosing among multiple AS-external-LSAs advertising the same destination. If you select Enable, the preference rules are those defined by RFC 1583. If you select Disable, the preference rules are those defined in Section 16.4.1 of the OSPF-2 standard (RFC 2328), which prevent routing loops when AS-external-LSAs for the same destination have been originated from different areas. The default value is Enable. To prevent routing loops, you should select Disable, but only if all OSPF routers in the routing domain are capable of operating according to RFC 2328.
ABR Status	The values of this are Enabled or Disabled. Enabled implies that the router is an area border router. Disabled implies that it is not an area border router.
Opaque LSA Status	Enables or disables the storing and flooding of opaque LSAs. An opaque LSA is used for flooding user-defined information within an OSPF router domain.
Exit Overflow Interval (sec)	Enter the number of seconds that, after entering overflow state, the router should wait before attempting to leave overflow state. This allows the router to again originate non-default AS-external-LSAs. If you enter 0, the router does not leave Overflow State until restarted. The range is 0 to 2147483647 seconds.
SPF DelayTime (secs)	Enter the number of seconds, Delay time (in seconds) is the time between when OSPF receives a topology change and when it starts an SPF calculation. It can be an integer from 0 to 65535. The default time is 5 seconds. A value of 0 means that there is no delay; that is, the SPF calculation is started immediately.
SPF HoldTime(secs)	Enter the number of seconds, minimum time (in seconds) between two consecutive SPF calculations.It can be an integer from 0 to 65535. The default time is 10 seconds. A value of 0 means that there is no delay; that is, two SPF calculations can be done, one immediately after the other.
External LSA Count	The number of external (LS type 5) LSAs (link state advertisements) in the link state database.
External LSA Checksum	The sum of the LS checksums of the external LSAs (link state advertisements) contained in the link-state database. This sum can be used to determine if there has been a change in a router's link state database, and to compare the link-state databases of two routers. This value is in hexadecimal.
AS_Opaque LSA Count	Number of type-11 Opaque Link State Advertisements (LSAs) received in the AS. Opaque LSAs consist of a standard LSA header followed by a 32-bit aligned application-specific information field. The information contained in Opaque LSAs may be used directly by OSPF or indirectly by some application wishing to distribute information throughout the OSPF domain. Link-state type 11 denotes that the LSA is flooded throughout the Autonomous System (AS).
AS_Opaque LSA Checksum	The sum of the checksums of type-11 Link Opaque LSA received in the AS. This sum can be used to determine if there has been a change in a router's link-state database, and to compare the link-state database of two routers.

Table 4-7: OSPF Configuration Fields (Continued)

Field	Description
New LSAs Originated	In any given OSPF area, a router originates several LSAs. Each router originates a router-LSA. If the router is also the Designated Router for any of the area's networks, it originates network-LSAs for those networks. This value represents the number of LSAs originated by this router.
LSAs Received	The number of LSAs (link state advertisements) received that were determined to be new instantiations. This number does not include newer instantiations of self-originated LSAs.
External LSDB Limit	The maximum number of AS-External-LSAs that can be stored in the database. A value of -1 implies there is no limit on the number that can be saved. The valid range of values is -1 to 2147483647.
Default Metric	Sets a default for the metric of redistributed routes.This field displays the default metric if one has already been set or blank if not configured earlier. The valid values are 1 to 16777214.
Maximum Paths	Configure the maximum number of paths that OSPF can report to a given destination. The valid values are 1 to 2.
Autocost Reference Bandwidth	This field configures the value that OSPF uses in calculating the default metric for an interface. OSPF calculates the link cost of each interface as: Cost = (Reference Bandwidth in Mbps) / (Interface Bandwidth). **Example**: Setting this value to 1000 Mbps would cause all 1-Gbps interfaces to have a default cost of 1000/1000 = 1. For 100 Mbps interfaces, the default cost would be 1000/100 = 10.
Default Passive Setting	Select whether OSPF interfaces default to passive mode. In passive mode, interfaces do not send OSPF routing updates. This setting is Disabled by default, so that all interfaces default to non-passive mode. If enabled, then all interfaces default to passive mode, and the network manager can selectively enable interfaces to send OSPF routing updates.
Default Route Advertise Fields	
Default Information Originate	Enable or Disable Default Route Advertise.
Always	Sets the router advertise 0.0.0.0/0.0.0.0 when set to True.
Metric	Specifies the metric of the default route. The valid values are (0 to 16777214)
Metric Type	Sets the metric type of the default route. Options are External Type 1 and External Type 2. External Type 2 is the default.

If you make changes to the page, click **Submit** to apply the changes to the system.

4.3.2 OSPF Area Configuration

The OSPF Area Configuration page lets you create a Stub area configuration and NSSA once you have enabled OSPF on an interface through **Routing > OSPF > Interface Configuration**. At least one router must have OSPF enabled for this web page to display.

To display the page, click **Routing > OSPF > Area Configuration** in the navigation menu.

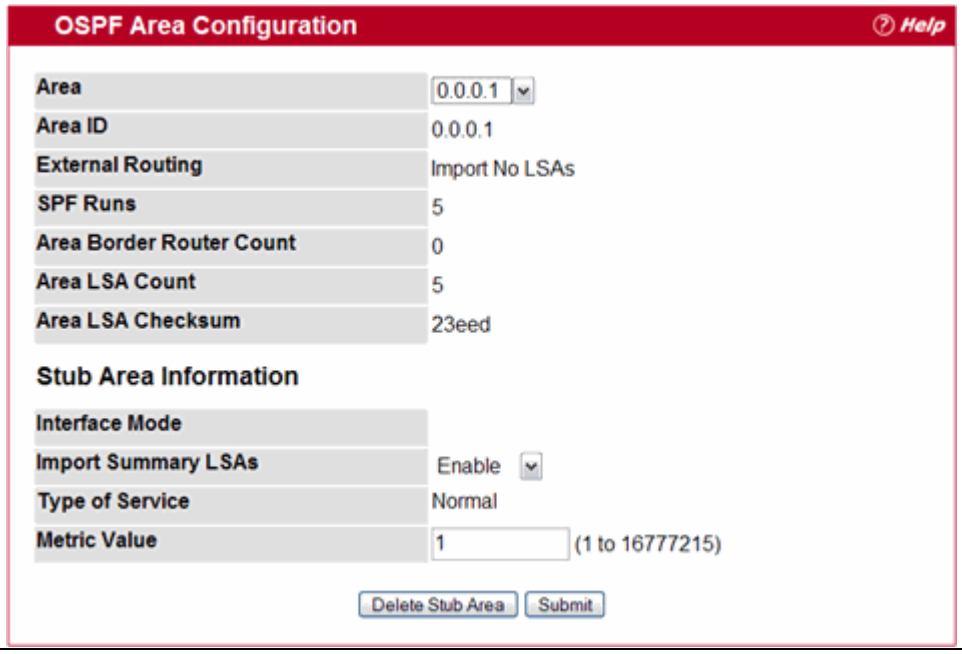

Figure 4-7: OSPF Area Configuration

Table 4-8: OSPF Area Configuration Fields

Field	Description
Area	Select the area to be displayed from the dropdown menu. When an area is selected, fields in the Stub Area Information are displayed.
Area ID	The OSPF area. An Area ID is a 32-bit integer in dotted decimal format that uniquely identifies the area to which a router interface connects.
External Routing	A definition of the router's capabilities for the area, including whether or not AS-external-LSAs are flooded into/throughout the area. If the area is a stub area, then these are the possible options for which you may configure the external routing capability, otherwise the only option is Import External LSAs.
SPF Runs	The number of times that the intra-area route table has been calculated using this area's link-state database. This is typically done using Dijkstra's algorithm.
Area Border Router Count	The total number of area border routers reachable within this area. This is initially zero, and is calculated in each SPF Pass.
Area LSA Count	The total number of link-state advertisements in this area's link-state database, excluding AS External LSAs.
Area LSA Checksum	The 32-bit unsigned sum of the link-state advertisements' LS checksums contained in this area's link-state database. This sum excludes external (LS type 5) link-state advertisements. The sum can be used to determine if there has been a change in a router's link state database, and to compare the link-state database of two routers. This value is in hexadecimal.
***Stub Area Information*—**The fields displayed depend on whether the area is a stub area or an NSSA. All possible fields are listed below.	
Interface Mode	This field tells you whether the area is or is not a stub area. If the area may be a stub area, a Create Stub Area button is displayed. If you have configured the area as a stub area a Delete Stub Area button is displayed. Otherwise neither button is displayed.
Import Summary LSAs	Select Enable or Disable from the dropdown menu. If you select Enable summary LSAs is imported into stub areas.

Table 4-8: OSPF Area Configuration Fields (Continued)

Field	Description
Type of Service	Select the type of metric specified in the Metric Value field: The type of metric for the stub area where valid types are: • **OSPF Metric**: Regular OSPF metric • **Comparable Cost**: External Type 1 metrics that are comparable to the OSPF metric • **Non-comparable Cost**: External Type 2 metrics that are assumed to be larger than the cost of the OSPF metric
Metric Value	Enter the metric value you want applied for the default route advertised into the stub area. Valid values range from 1 to 16,777,215.
Translator Role	Configure the NSSA Translator Role as always/candidate. This field only displays if the area is an NSSA area.
Translator Stability Interval	Configure the Translator Stability Interval for the selected NSSA. This field only displays if the area is an NSSA area.
No-Redistribute Mode	Configure the route redistribution for the selected NSSA. This field only displays if the area is an NSSA area.
Translator State	Displays the state of the Translator. This field only displays if the area is an NSSA area.

If you make any changes to the page, click **Submit** to apply the changes to the system.

4.3.3 OSPF Stub Area Summary

Use the OSPF Stub Area Summary page to display OSPF stub area detail.

To display the page, click **Routing > OSPF > Stub Area Summary** in the navigation tree.

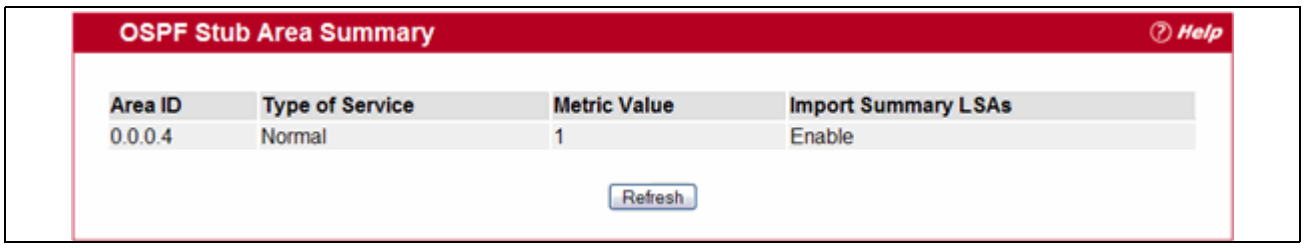

Figure 4-8: OSPF Stub Area Summary

Table 4-9: OSPF Stub Area Summary Fields

Field	Description
Area ID	The Area ID of the Stub area.
Type of Service	The type of service associated with the stub metric. The switch supports Normal only.
Metric Value	Displays the configured metric value.
Import Summary LSAs	Displays whether the import of Summary LSAs is Enabled or Disabled.

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

4.3.4 OSPF Area Range Configuration

Use the OSPF Area Range Configuration page to configure and display an area range for a specified NSSA.

To display the page, click **Routing > OSPF > Area Range Configuration** in the navigation menu.

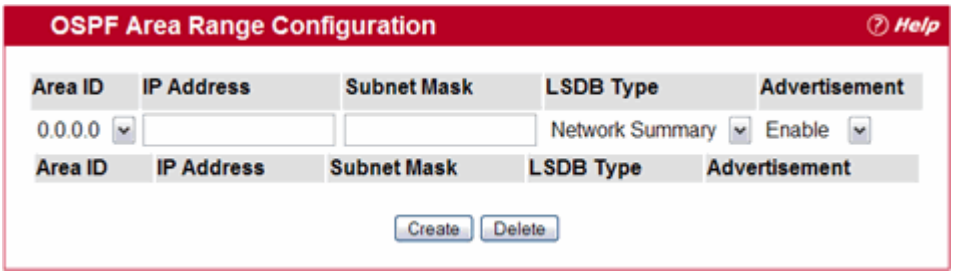

Figure 4-9: OSPF Area Range Configuration

Table 4-10: OSPF Area Range Configuration Fields

Field	Description
Area ID	Select the area for which data is to be configured from the dropdown menu.
IP Address	Enter the IP Address for the address range for the selected area.
Subnet Mask	Enter the Subnet Mask for the address range for the selected area.
LSDB Type	Select the type of Link Advertisement associated with the specified area and address range. The default type is 'Network Summary.
Advertisement	Select Enable or Disable from the dropdown menu. If you selected Enable the address range is advertised outside the area via a Network Summary LSA. The default is Enable.
OSPF Area Range Table	
Area ID	Displays the OSPF area.
IP Address	Displays the IP address of an address range for the area.
Subnet Mask	Displays the subnet mask of an address range for the area.
LSDB Type	Displays the link advertisement type for the address range and area.
Advertisement	Displays the advertisement mode for the address range and area.

- To configure an OSPF area range, select an ID from the Area ID dropdown menu, enter or select the desired parameters from the configurable fields, and then click **Create**. After you click Create, the page refreshes, and the new OSPF area range appears in the table.
- To delete an existing range, select the area ID from the dropdown menu, enter the IP address and subnet mask in the appropriate fields, and then click **Delete**.

4.3.5 OSPF Interface Statistics

Use the OSPF Interface Statistics page to display statistics for the selected interface. The information is displayed only if OSPF is enabled.

To display the page, click **Routing > OSPF > Interface Statistics** in the navigation tree.

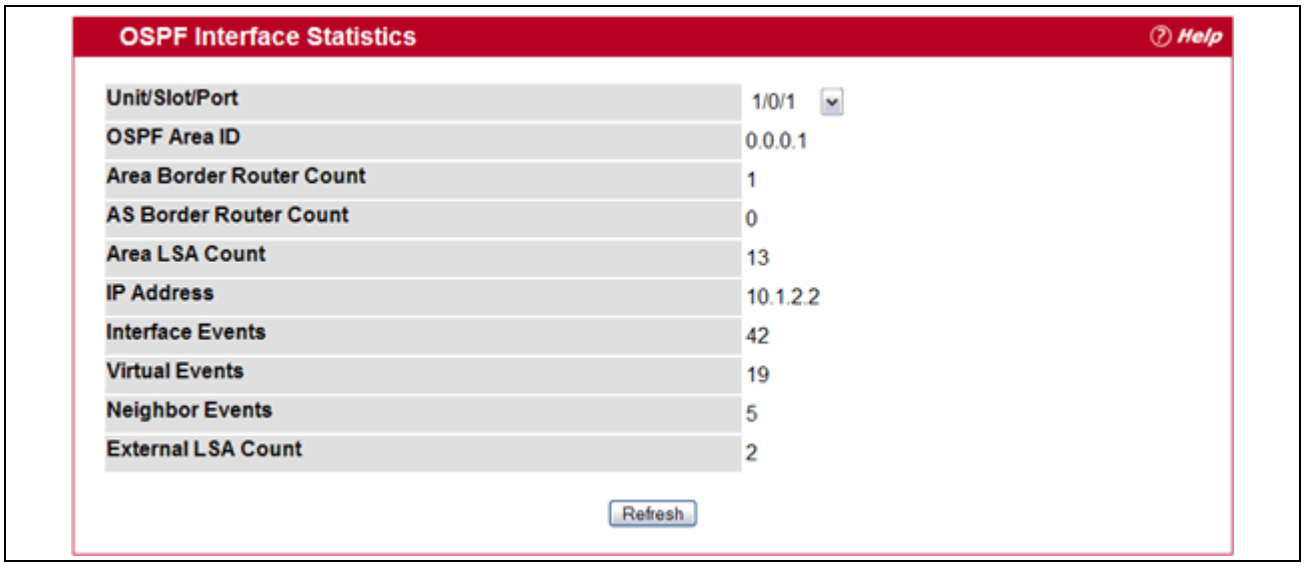

Figure 4-10: OSPF Interface Statistics

Table 4-11: OSPF Interface Statistics Fields

Field	Description
Interface	Select the interface for which data is to be displayed from the dropdown menu.
OSPF Area ID	The OSPF area to which the selected router interface belongs. An OSPF Area ID is a 32-bit integer in dotted decimal format that uniquely identifies the area to which the interface connects.
Area Border Router Count	The total number of area border routers reachable within this area. This is initially zero, and is calculated in each SPF Pass.
AS Border Router Count	The total number of Autonomous System border routers reachable within this area. This is initially zero, and is calculated in each SPF Pass.
Area LSA Count	The total number type-10 link-state advertisements in this area's link-state database, excluding AS External LSAs. Link-state type 10 denotes an area-local scope. Type-10 Opaque LSAs are not flooded beyond the borders of their associated area.
IP Address	The IP address of the interface.
Interface Events	The number of times the specified OSPF interface has changed its state, or an error has occurred.
Virtual Events	The number of state changes or errors that have occurred on this virtual link.
Neighbor Events	The number of times this neighbor relationship has changed state, or an error has occurred.
External LSA Count	The number of external (LS type 5) link-state advertisements in the link-state database.

4.3.6 OSPF Interface Configuration

Use the OSPF Interface Configuration page to configure an OSPF interface.

To display the page, click **Routing > OSPF > Interface Configuration** in the navigation tree.

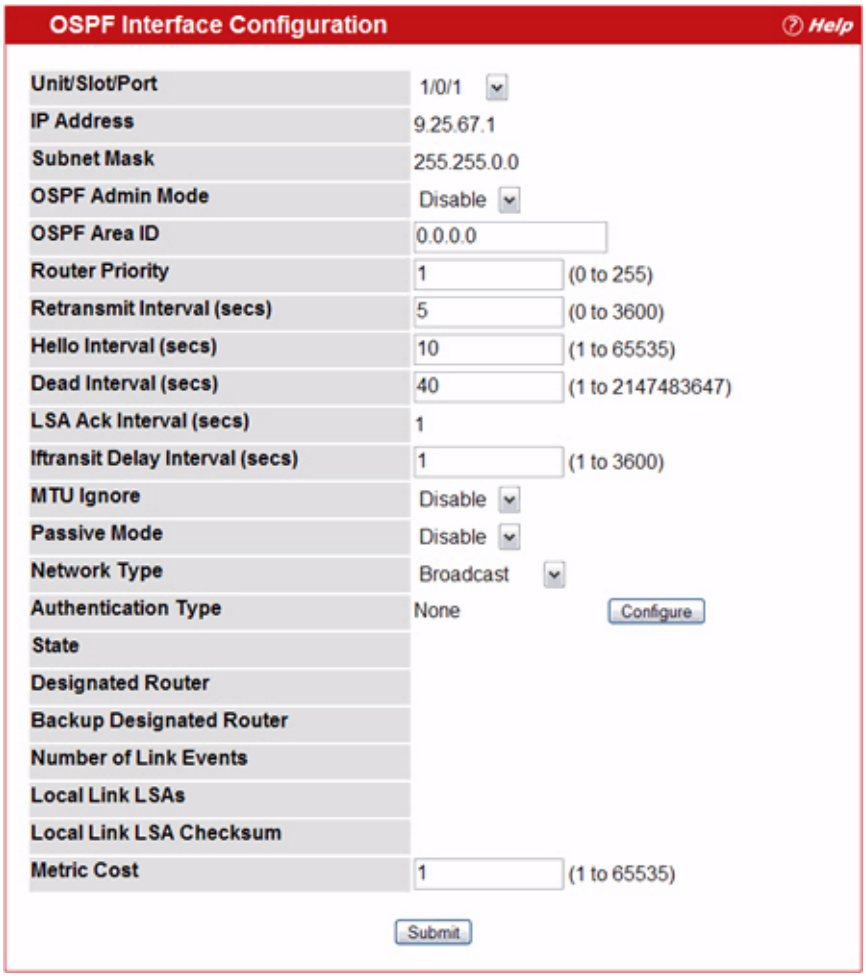

Figure 4-11: OSPF Interface Configuration

Table 4-12: OSPF Interface Configuration Fields

Field	Description
Slot/Port	Select the interface for which data is to be displayed or configured from the dropdown menu.
IP Address	Displays the address of the VLAN Interface.
Subnet Mask	Displays the subnet mask of the VLAN Interface.
OSPF Admin Mode	You may select Enable or Disable from the dropdown menu. The default value is **Disable**. You can configure OSPF parameters without enabling OSPF Admin Mode, but they have no effect until Admin Mode is enabled. The following information is displayed only if the Admin Mode is enabled: State, Designated Router, Backup Designated Router, Number of Link Events, LSA Ack Interval, and Metric Cost. For OSPF to be fully functional, you must enter a valid IP Address and Subnet Mask via the Interface IP Configuration page or through the CLI command: `config ip interface network`. **Note**: Once OSPF is initialized on the router, it remains initialized until the router is reset.
OSPF Area ID	Enter the 32-bit integer in dotted decimal format that uniquely identifies the OSPF area to which the selected router interface connects. If you assign an Area ID which does not exist, the area is created with default values.

Table 4-12: OSPF Interface Configuration Fields (Continued)

Field	Description
Router Priority	Enter the OSPF priority for the selected interface. The priority of an interface is specified as an integer from 0 to 255. The default is 1, which is the highest router priority. A value of 0 indicates that the router is not eligible to become the designated router on this network.
Retransmit Interval (secs)	Enter the OSPF retransmit interval for the specified interface. This is the number of seconds between link-state advertisements for adjacencies belonging to this router interface. This value is also used when retransmitting database descriptions and link-state request packets. Valid values range from 1 to 3600 seconds (1 hour). The default is 5 seconds.
Hello Interval (secs)	Enter the OSPF hello interval for the specified interface in seconds. This parameter must be the same for all routers attached to a network. Valid values range from 1 to 65,535. The default is 10 seconds.
Dead Interval (secs)	Enter the OSPF dead interval for the specified interface in seconds. This specifies how long a router waits to see a neighbor router's Hello packets before declaring that the router is down. This parameter must be the same for all routers attached to a network. This value should a multiple of the Hello Interval (e.g., 4). Valid values range from 1 to 2147483647. The default is 40.
LSA Ack Interval	The number of seconds between LSA Acknowledgment packet transmissions, which must be less than the Retransmit Interval.
Iftransit Delay Interval (secs)	Enter the OSPF Transit Delay for the specified interface. This specifies the estimated number of seconds it takes to transmit a link state update packet over the selected interface. Valid values range from 1 to 3600 seconds (1 hour). The default value is 1 second.
MTU Ignore	Disables OSPF MTU mismatch detection on receiving packets. The default value is Disable.
Passive Mode	When you enable passive mode on an OSPF interface, you disable sending OSPF routing updates on the interface. An OSPF adjacency will not be formed on a passive interface. Subnet prefixes for IP addresses configured on the interface will continue to be advertised as stub networks.
Authentication Type	You may select an authentication type other than None by clicking the **Configure** button. You then see a new web page, where you can select the authentication type from the dropdown menu. Possible values are: • **None**: This is the initial interface state. If you select this option from the dropdown menu on the second screen you are returned to the first screen, and no authentication protocols are run. • **Simple**: If you select Simple, you are prompted to enter an authentication key. This key is included, in the clear, in the OSPF header of all packets sent on the network. All routers on the network must be configured with the same key. • **Encrypt**: If you select Encrypt, you are prompted to enter both an authentication key and an authentication ID. Encryption uses the MD5 Message-Digest algorithm. All routers on the network must be configured with the same key and ID.
Interface Type	The OSPF interface type, which is always broadcast.

Table 4-12: OSPF Interface Configuration Fields (Continued)

Field	Description
State	The current state of the selected router interface. Possible values are: • **Down**: This is the initial interface state. In this state, the lower-level protocols have indicated that the interface is unusable. In this state, interface parameters are set to their initial values. All interface timers are disabled, and there are no adjacencies associated with the interface. • **Loopback**: In this state, the router's interface to the network is looped back either in hardware or software. The interface is unavailable for regular data traffic. However, it may still be desirable to gain information on the quality of this interface, either through sending ICMP pings to the interface or through something like a bit error test. For this reason, IP packets may still be addressed to an interface in Loopback state. To facilitate this, such interfaces are advertised in router-LSAs as single host routes, whose destination is the IP interface address. • **Waiting**: The router is trying to determine the identity of the (Backup) Designated Router for the network by monitoring received Hello Packets. The router is not allowed to elect a Backup Designated Router or a Designated Router until it transitions out of Waiting state. This prevents unnecessary changes of (Backup) Designated Router. • **Designated Router**: This router is itself the Designated Router on the attached network. Adjacencies are established to all other routers attached to the network. The router must also originate a network-LSA for the network node. The network-LSA contains links to all routers (including the Designated Router itself) attached to the network. • **Backup Designated Router**: This router is itself the Backup Designated Router on the attached network. It is promoted to Designated Router if the present Designated Router fails. The router establishes adjacencies to all other routers attached to the network. The Backup Designated Router performs slightly different functions during the Flooding Procedure, as compared to the Designated Router. • **Other Designated Router**: The interface is connected to a broadcast or NBMA network on which other routers have been selected to be the Designated Router and Backup Designated Router either. The router attempts to form adjacencies to both the Designated Router and the Backup Designated Router. The State is only displayed if the OSPF admin mode is enabled.
Designated Router	The identity of the Designated Router for this network, in the view of the advertising router. The Designated Router is identified here by its router ID. The value 0.0.0.0 means that there is no Designated Router. This field is only displayed if the OSPF admin mode is enabled.
Backup Designated Router	The identity of the Backup Designated Router for this network, in the view of the advertising router. The Backup Designated Router is identified here by its router ID. Set to 0.0.0.0 if there is no Backup Designated Router. This field is only displayed if the OSPF admin mode is enabled.
Number of Link Events	This is the number of times the specified OSPF interface has changed its state. This field is only displayed if the OSPF admin mode is enabled.
Local Link LSAs	Number of Opaque Link State Advertisements (type-9) received on the interface. Opaque LSAs provide a generalized mechanism to allow for the future extensibility of OSPF. The information contained in Opaque LSAs may be used directly by OSPF or indirectly by some application wishing to distribute information throughout the OSPF domain. Opaque LSAs consist of a standard LSA header followed by a 32-bit aligned application-specific information field. Link-state type 9 denotes a link-local scope. Type-9 Opaque LSAs are not flooded beyond the local network.
Local Link LSA Checksum	The sum of the Local Link Opaque LSA checksums received on the interface. This sum can be used to determine if there has been a change in a router's link-state database, and to compare the link-state database of two routers.
Metric Cost	Enter the value on this interface for the cost. The range for the metric cost is between 1 and 65,535. Metric Cost is only configurable/displayed if OSPF is initialized on the interface.

4.3.6.1 Configuring an OSPF Interface Configuration

1. From the OSPF Interface Configuration page, specify an interface to configure.
2. Specify values in the remaining fields as needed.
3. To configure the authentication, click **Configure**.

   The page refreshes and displays the OSPF Interface Authentication Configuration page.

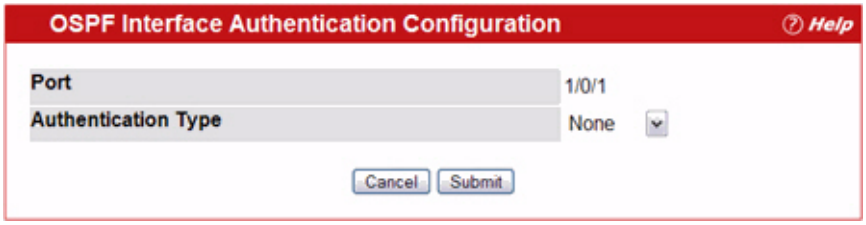

Figure 4-12: OSPF Interface Authentication Configuration

4. Select the type of authentication to use.

   If you select Simple or Encrypt as the authentication, the screen refreshes, and additional fields display. Enter the required information into the new fields.
5. Click **Submit** to apply the changes to the system and return to the OSPF Interface Configuration page.
6. To cancel the authentication configuration and return to the OSPF Interface Configuration page, click **Cancel**.

   The OSPF interface is configured.

4.3.7 Neighbor Table

Use the OSPF Neighbor Table page to display the OSPF neighbor table list. When a particular neighbor ID is specified, detailed information about a neighbor is given. The information below is only displayed if OSPF is enabled.

To display the page, click **Routing > OSPF > Neighbor Table** in the navigation tree.

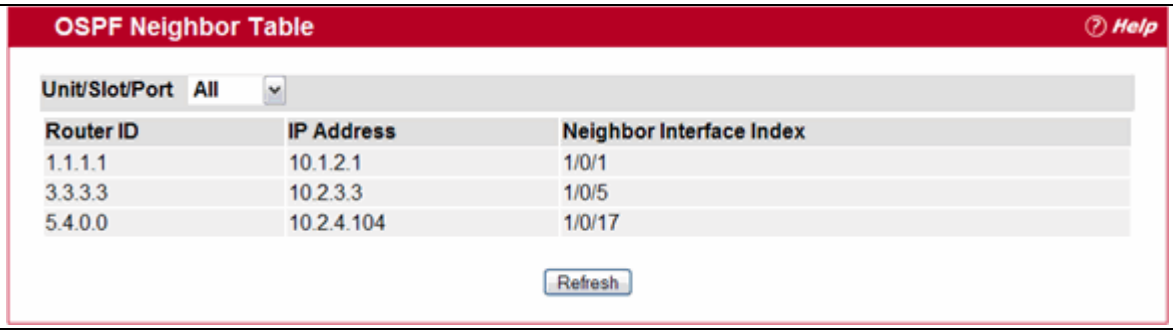

Figure 4-13: OSPF Neighbor Table

Table 4-13: OSPF Neighbor Table Fields

Field	Description
Slot/Port	Select the interface for which data is to be displayed from a dropdown menu.
Router ID	A 32-bit integer in dotted decimal format representing the neighbor interface.
IP Address	The IP address of the neighboring router's interface to the attached network. It is used as the destination IP address when protocol packets are sent as unicasts along this adjacency. Also used in router-LSAs as the Link ID for the attached network if the neighboring router is selected to be designated router. The Neighbor IP address is learned when Hello packets are received from the neighbor. For virtual links, the Neighbor IP address is learned during the routing table build process.
Neighbor Interface Index	An interface identifying the neighbor interface index.

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

4.3.8 OSPF Neighbor Configuration

Use the OSPF Neighbor Configuration page to display the OSPF neighbor configuration for a selected neighbor ID. When a particular neighbor ID is specified, detailed information about a neighbor is given. The information below is only displayed if OSPF is enabled and the interface has a neighbor. The IP address is the IP address of the neighbor.

To display the page, click **Routing > OSPF > Neighbor Configuration** in the navigation tree.

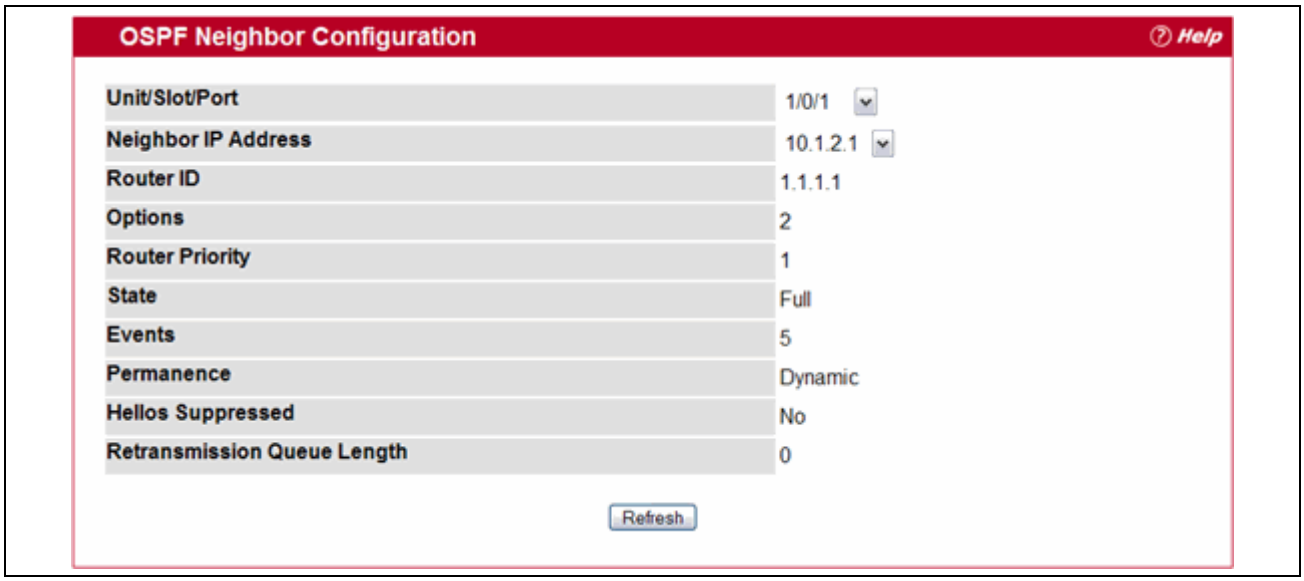

Figure 4-14: OSPF Neighbor Configuration

Table 4-14: OSPF Neighbor Configuration Fields

Field	Description
VLAN	Select the VLAN interface on which routing is enabled. FASTPATH also supports port-based routing interfaces.
Neighbor IP Address	Select the IP Address of the neighbor for which data is to be displayed.
Router ID	A 32-bit integer in dotted decimal format that identifies the neighbor router.
Options	The optional OSPF capabilities supported by the neighbor. The OSPF Options field is present in OSPF Hello packets, Database Description packets, and all link-state advertisements. The Options field enables OSPF routers to support (or not support) optional capabilities, and to communicate their capability level to other OSPF routers. Through this mechanism, routers of differing capabilities can be mixed within an OSPF routing domain. The Options value is a bitmap, and it signifies the capability of the neighbor. The option field is defined as follows: • E-bit: Describes the way AS-external-LSAs are flooded as described in Sections 3.6, 9.5, 10.8 and 12.1.2 of RFC 2328, “OSPF Version 2”. • MC-bit: Describes whether IP multicast datagrams are forwarded according to the specifications in RFC 1584, “Multicast Extensions to OSPF”. • N/P-bit: Describes the handling of Type-7 LSAs, as specified in RFC 1587, "The OSPF NSSA Option". • EA-bit: This bit describes the router's willingness to receive and forward External-Attributes-LSAs, as specified in "The OSPF External Attributes LSA". • DC-bit: This bit describes the router's handling of demand circuits, as specified in RFC 1793 "Extending OSPF to Support Demand Circuits".
Router Priority	Displays the OSPF priority for the specified neighbor. The priority of a neighbor is a priority integer from 0 to 255. A value of 0 indicates that the router is not eligible to become the designated router on this network.
State	The state of a neighbor can be the following: • **Down**: This is the initial state of a neighbor conversation. It indicates that there has been no recent information received from the neighbor. On NBMA networks, Hello packets may still be sent to Down neighbors, although at a reduced frequency. • **Init**: In this state, a Hello packet has recently been seen from the neighbor. However, bidirectional communication has not yet been established with the neighbor (i.e., the router itself did not appear in the neighbor's Hello packet). All neighbors in this state (or greater) are listed in the Hello packets sent from the associated interface. • **2-Way**: In this state, communication between the two routers is bidirectional. This has been assured by the operation of the Hello Protocol. This is the most advanced state short of beginning adjacency establishment. The (Backup) Designated Router is selected from the set of neighbors in state 2-Way or greater. • **Exchange Start**: This is the first step in creating an adjacency between the two neighboring routers. The goal of this step is to decide which router is the master, and to decide upon the initial DD sequence number. Neighbor conversations in this state or greater are called adjacencies. • **Exchange**: In this state, the router is describing its entire link state database by sending Database Description packets to the neighbor. In this state, Link State Request Packets may also be sent asking for the neighbor's more recent LSAs. All adjacencies in Exchange state or greater are used by the flooding procedure. These adjacencies are fully capable of transmitting and receiving all types of OSPF routing protocol packets. • **Loading**: In this state, Link State Request packets are sent to the neighbor asking for the more recent LSAs that have been discovered (but not yet received) in the Exchange state. • **Full**: In this state, the neighboring routers are fully adjacent. These adjacencies appear in router-LSAs and network-LSAs.

Table 4-14: OSPF Neighbor Configuration Fields (Continued)

Field	Description
Events	The number of times this neighbor relationship has changed state, or an error has occurred.
Permanence	This variable displays the status of the entry. Dynamic and permanent refer to how the neighbor became known.
Hellos Suppressed	This indicates whether Hellos are being suppressed to the neighbor.
Retransmission Queue Length	The current length of the retransmission queue.

4.3.9 OSPF Link State Database

Use the OSPF Link State Database page to display OSPF link state information.

To display the page, click **Routing > OSPF > Link State Database** in the navigation tree.

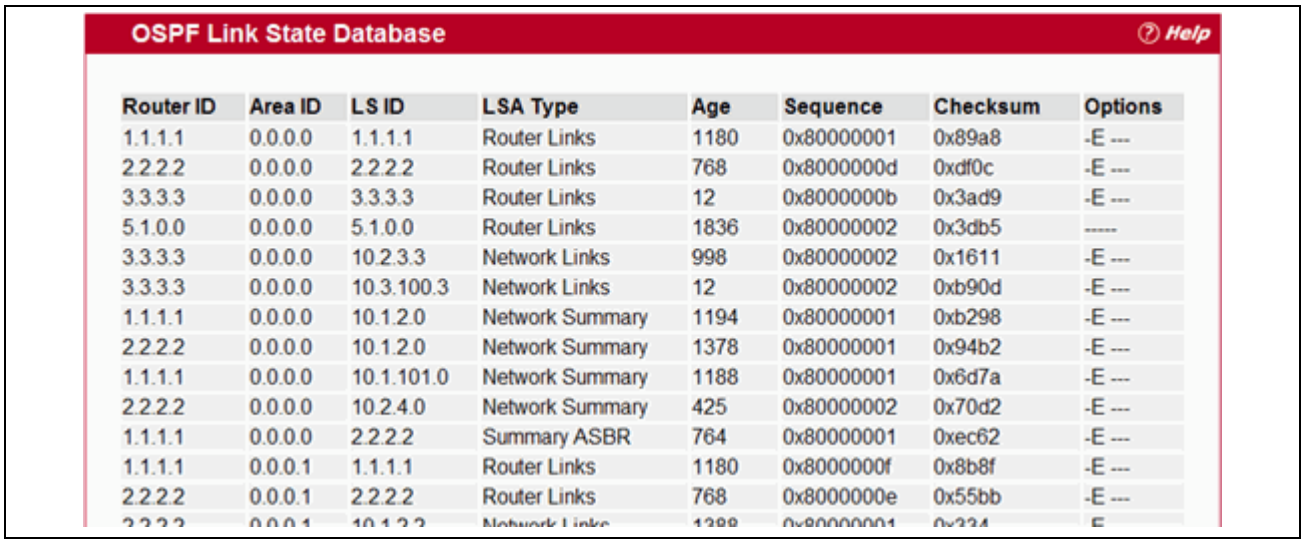

OSPF Link State Database

Help

Router ID	Area ID	LS ID	LSA Type	Age	Sequence	Checksum	Options
1.1.1.1	0.0.0.0	1.1.1.1	Router Links	1180	0x80000001	0x89a8	-E ---
2.2.2.2	0.0.0.0	2.2.2.2	Router Links	768	0x8000000d	0xdf0c	-E ---
3.3.3.3	0.0.0.0	3.3.3.3	Router Links	12	0x8000000b	0x3ad9	-E ---
5.1.0.0	0.0.0.0	5.1.0.0	Router Links	1836	0x80000002	0x3db5	-----
3.3.3.3	0.0.0.0	10.2.3.3	Network Links	998	0x80000002	0x1611	-E ---
3.3.3.3	0.0.0.0	10.3.100.3	Network Links	12	0x80000002	0xb90d	-E ---
1.1.1.1	0.0.0.0	10.1.2.0	Network Summary	1194	0x80000001	0xb298	-E ---
2.2.2.2	0.0.0.0	10.1.2.0	Network Summary	1378	0x80000001	0x94b2	-E ---
1.1.1.1	0.0.0.0	10.1.101.0	Network Summary	1188	0x80000001	0x6d7a	-E ---
2.2.2.2	0.0.0.0	10.2.4.0	Network Summary	425	0x80000002	0x70d2	-E ---
1.1.1.1	0.0.0.0	2.2.2.2	Summary ASBR	764	0x80000001	0xec62	-E ---
1.1.1.1	0.0.0.1	1.1.1.1	Router Links	1180	0x8000000f	0x8b8f	-E ---
2.2.2.2	0.0.0.1	2.2.2.2	Router Links	768	0x8000000e	0x55bb	-E ---

Figure 4-15: OSPF Link State Database

Table 4-15: OSPF Link State Database Fields

Field	Description
Router ID	The 32-bit integer in dotted decimal format that uniquely identifies the router within the autonomous system (AS). The Router ID is set on the IP Configuration page. If you want to change the Router ID you must first disable OSPF. After you set the new Router ID, you must re-enable OSPF to have the change take effect. The default value is 0.0.0.0, although this is not a valid Router ID.
Area ID	The ID of an OSPF area to which one of the router interfaces is connected. An Area ID is a 32-bit integer in dotted decimal format that uniquely identifies the area to which an interface is connected.
LSA Type	The format and function of the link state advertisement. Possible values are: • Router Links • Network Links • Network Summary • ASBR Summary • AS-external
LS ID	The Link State ID identifies the piece of the routing domain that is being described by the advertisement. The value of the LS ID depends on the advertisement's LS type.
Age	The time since the link state advertisement was first originated, in seconds.
Sequence	The sequence number field is a signed 32-bit integer. It is used to detect old and duplicate link state advertisements. The larger the sequence number, the more recent the advertisement.
Checksum	The checksum is used to detect data corruption of an advertisement. This corruption can occur while an advertisement is being flooded, or while it is being held in a router's memory. This field is the checksum of the complete contents of the advertisement, except the LS age field.
Options	The Options field in the link state advertisement header indicates which optional capabilities are associated with the advertisement. Possible values are: **Q**: This enables support for QoS Traffic Engineering. **E**: This describes the way AS-external-LSAs are flooded. **MC**: This describes the way IP multicast datagrams are forwarded according to the standard specifications. **O**: This describes whether Opaque-LSAs are supported. **V**: This describes whether OSPF++ extensions for VPN/COS are supported.

4.3.10 OSPF Virtual Link Configuration

Use the OSPF Virtual Link Configuration page to create or configure virtual interface information for a specific area and neighbor. A valid OSPF area must be configured before this page can be displayed.

To display the page, click **Routing > OSPF > Virtual Link Configuration** in the navigation tree.

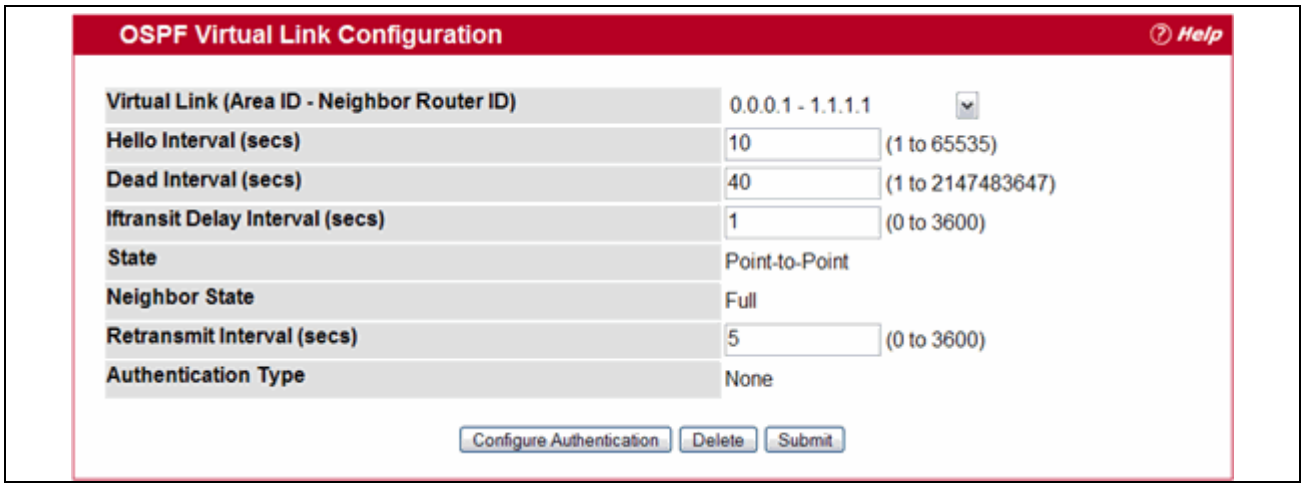

Figure 4-16: OSPF Virtual Link Configuration

Table 4-16: OSPF Virtual Link Configuration Fields

Field	Description
Create New Virtual Link	Select this option from the menu to define a new virtual link. The area portion of the virtual link identification is fixed: you are prompted to enter the Neighbor Router ID on a new screen.
Area ID	**Neighbor Router ID:** Select the virtual link for which you want to display or configure data. It consists of the Area ID and Neighbor Router ID.
Hello Interval	Enter the OSPF hello interval for the specified interface in seconds. This parameter must be the same for all routers attached to a network. Valid values range from 1 to 65,535. The default is 10 seconds.
Dead Interval	Enter the OSPF dead interval for the specified interface in seconds. This specifies how long a router waits to see a neighbor router's Hello packets before declaring that the router is down. This parameter must be the same for all routers attached to a network. This value should a multiple of the Hello Interval (e.g., 4). Valid values range from 1 to 2147483647. The default is 40.
Iftransit Delay Interval (secs)	The OSPF Transit Delay for the virtual link in units of seconds. It specifies the estimated number of seconds it takes to transmit a link state update packet over this interface.

Table 4-16: OSPF Virtual Link Configuration Fields (Continued)

Field	Description
State	The current state of the selected Virtual Link. One of: • **Down**: This is the initial interface state. In this state, the lower-level protocols have indicated that the interface is unusable. In this state, interface parameters are set to their initial values. All interface timers are disabled, and there are no adjacencies associated with the interface. • **Waiting**: The router is trying to determine the identity of the (Backup) Designated Router by monitoring received Hello Packets. The router is not allowed to elect a Backup Designated Router or a Designated Router until it transitions out of Waiting state. This prevents unnecessary changes of (Backup) Designated Router. • **Point-to-Point**: The interface is operational, and is connected either to the virtual link. On entering this state the router attempts to form an adjacency with the neighboring router. Hello Packets are sent to the neighbor every HelloInterval seconds. • **Designated Router**: This router is itself the Designated Router on the attached network. Adjacencies are established to all other routers attached to the network. The router must also originate a network-LSA for the network node. The network-LSA contains links to all routers (including the Designated Router itself) attached to the network. • **Backup Designated Router**: This router is itself the Backup Designated Router on the attached network. It is promoted to Designated Router if the present Designated Router fails. The router establishes adjacencies to all other routers attached to the network. The Backup Designated Router performs slightly different functions during the Flooding Procedure, as compared to the Designated Router. • **Other Designated Router**: The interface is connected to a broadcast or NBMA network on which other routers have been selected to be the Designated Router and Backup Designated Router either. The router attempts to form adjacencies to both the Designated Router and the Backup Designated Router.
Neighbor State	The state of the Virtual Neighbor Relationship.
Retransmit Interval	Enter the OSPF retransmit interval for the specified interface. This is the number of seconds between link-state advertisements for adjacencies belonging to this router interface. This value is also used when retransmitting database descriptions and link-state request packets. Valid values range from 1 to 3600 seconds (1 hour). The default is 5 seconds.
Authentication Type	You may select an authentication type other than none by clicking on the Configure Authentication button. You then see a new screen, where you can select the authentication type from the dropdown menu. The choices are: • **None**: This is the initial interface state. If you select this option from the dropdown menu on the second screen you are returned to the first screen. • **Simple**: If you select Simple you are prompted to enter an authentication key. This key is included, in the clear, in the OSPF header of all packets sent on the network. All routers on the network must be configured with the same key. • **Encrypt**: If you select Encrypt you are prompted to enter both an authentication key and an authentication ID. Encryption uses the MD5 Message-Digest algorithm. All routers on the network must be configured with the same key and ID.
Authentication Key	Enter the OSPF Authentication Key for the specified interface. If you do not choose to use authentication you are not prompted to enter a key. If you choose Simple authentication you cannot use a key of more than 8 octets. If you choose Encrypt the key may be up to 16 octets long. The key value is only displayed if you are logged on with Read/Write privileges, otherwise it is displayed as asterisks.
Authentication ID	Enter the ID to be used for authentication. You are only prompted to enter an ID when you select Encrypt as the authentication type. The ID is a number between 0 ad 255, inclusive.

4.3.11 OSPF Virtual Link Summary

Use the OSPF Virtual Link Summary page to display all of the configured virtual links.

To display the page, click **Routing > OSPF > Virtual Link Summary** in the navigation tree.

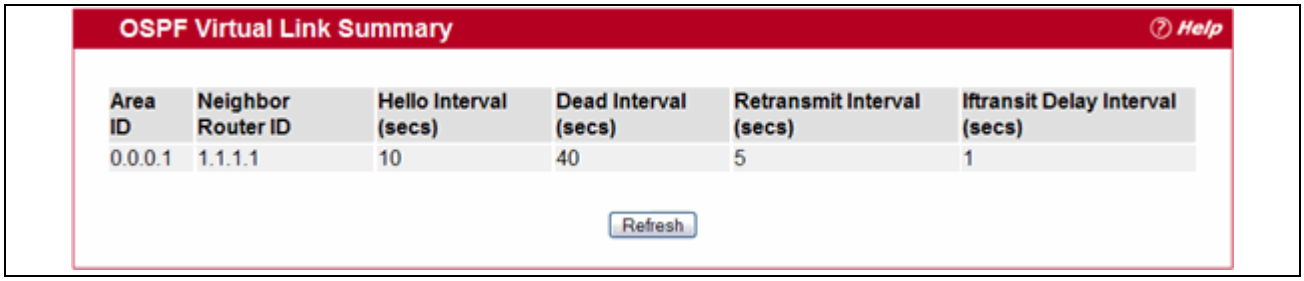

Figure 4-17: OSPF Virtual Link Summary

Table 4-17: OSPF Virtual Link Summary Fields

Field	Description
Area ID	The Area ID portion of the virtual link identification for which data is to be displayed. The Area ID and Neighbor Router ID together define a virtual link.
Neighbor Router ID	The neighbor portion of the virtual link identification. Virtual links may be configured between any pair of area border routers having interfaces to a common (non-backbone) area.
Hello Interval (secs)	The OSPF hello interval for the virtual link in units of seconds. The value for hello interval must be the same for all routers attached to a network.
Dead Interval (secs)	The OSPF dead interval for the virtual link in units of seconds. This specifies how long a router waits to see a neighbor router's Hello packets before declaring that the router is down. This parameter must be the same for all routers attached to a common network, and should be a multiple of the Hello Interval (i.e., 4).
Retransmit Interval (secs)	The OSPF retransmit interval for the virtual link in units of seconds. This specifies the time between link-state advertisements for adjacencies belonging to this router interface. This value is also used when retransmitting database descriptions and link-state request packets.
Iftransit Delay Interval (secs)	The OSPF Transit Delay for the virtual link in units of seconds. It specifies the estimated number of seconds it takes to transmit a link state update packet over this interface.

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

4.3.12 OSPF Route Redistribution Configuration

Use the OSPF Route Redistribution Configuration page to configure which routes are redistributed to other routes using OSPF depending on how they were learned—through Static configuration, directly connected hosts, RIP, and BGP. You can choose to redistribute routes learned from all of them or from selected sources.

To display the page, click **Routing > OSPF > Route Redistribution Configuration** in the navigation tree.

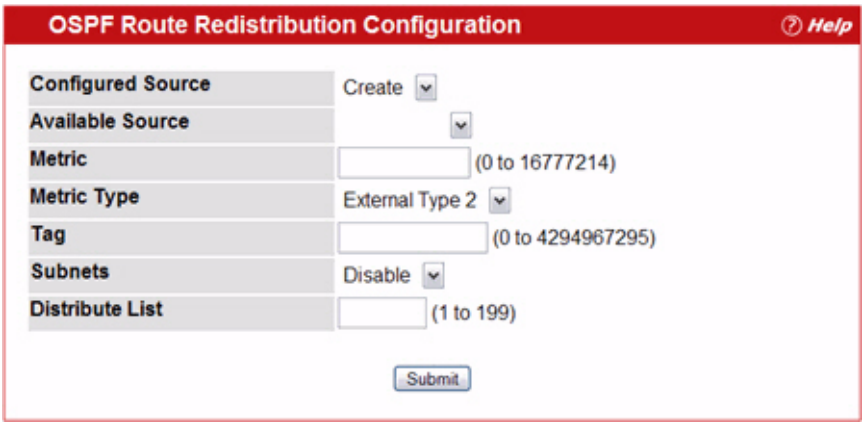

Figure 4-18: OSPF Route Redistribution Configuration

Table 4-18: OSPF Route Redistribution Configuration Fields

Field	Description
Configured Source	A protocol configured for OSPF to redistribute the routes learned through this protocol. Only source routes that have been configured for redistribute by OSPF are available. Create allows you to configure a new source route.
Available Source	A protocol available for configuration for OSPF to redistribute the routes. This box appears only if you select Create as Configured Source. Possible values are: • **Static**: The route was manually configured. • Connected: The route was determined automatically because the host is directly connected. • **RIP**: The route was determined through RIP. • **BGP**: The route was determined through BGP.
Metric	Sets the metric value for redistributed routes. This field displays a metric value if the source was preconfigured. The valid values are 0 to 16777214.
Metric Type	Select the OSPF metric type of redistributed routes from the dropdown menu.
Tag	Sets the tag field in routes redistributed. This field displays a tag value if the source was preconfigured, otherwise 0 is displayed. The valid values are 0 to 4294967295.
Subnets	Select whether the subnetted routes should be redistributed or not from the dropdown menu.
Distribute List	Selects the Access List that filters the routes to be redistributed by the destination protocol. Only permitted routes are redistributed. If this command refers to a non-existent access list, all routes are permitted. The dropdown menu lists the ACLs configured from the **Switching > Network Security > Access Control Lists > IP Access Control Lists** pages. When used for route filtering, the only fields in an access list that get used are: • Source IP Address and netmask • Destination IP Address and netmask • Action (permit or deny)

All other fields (source and destination port, precedence, tos, etc.) are ignored.

The source IP address is compared to the destination IP address of the route. The source IP netmask in the access list rule is treated as a wildcard mask, indicating which bits in the source IP address must match the destination address of the route. (Note that a 1 in the mask indicates a Don't Care in the corresponding address bit.)

When an access list rule includes a destination IP address and netmask (an extended access list), the destination IP address is compared to the network mask of the route destination. The destination netmask in the access list serves as a wildcard mask, indicating which bits in the route's destination mask are significant for the filtering operation.

If you make any changes to the page, click **Submit** to apply the changes to the system.

4.3.13 OSPF Route Redistribution Summary

Use the OSPF Route Redistribution Summary page to display OSPF Route Redistribution configurations.

To display the page, click **Routing > OSPF > Route Redistribution Summary** in the navigation tree.

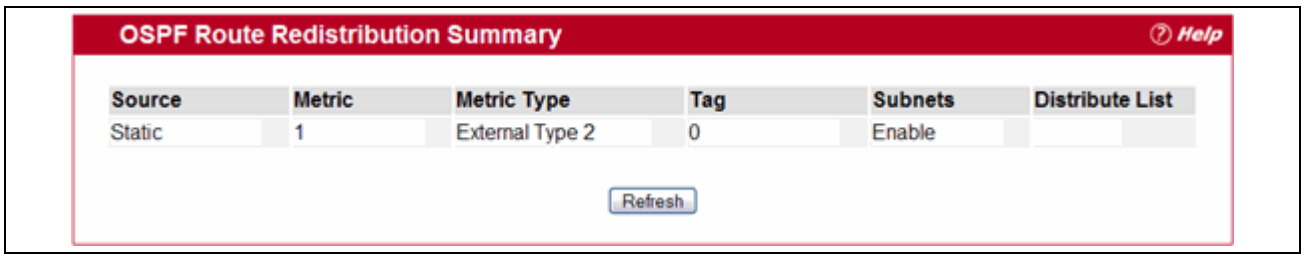

Figure 4-19: OSPF Route Redistribution Summary

Table 4-19: OSPF Route Redistribution Summary Fields

Field	Description
Source	The Source Route to be Redistributed by OSPF.
Metric	The Metric of redistributed routes for the given Source Route. Displays Unconfigured when not configured.
Metric Type	The OSPF metric type of redistributed routes.
Tag	The tag field in routes redistributed. This field displays the tag value if the source was preconfigured, otherwise 0 is displayed.
Subnets	Specifies whether the subnetted routes should be redistributed or not.
Distribute List	The Access List that filters the routes to be redistributed by the Destination Protocol.

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

4.4 Managing the BOOTP/DHCP Relay Agent

BootP/DHCP Relay Agent enables BootP/DHCP clients and servers to exchange BootP/DHCP messages across different subnets. The relay agent receives the requests from the clients, and checks the valid hops and giaddr fields. If the number of hops is greater than the configured, the agent assumes the packet is looped through the agents and discards the packet. If giaddr field is zero the agent must fill in this field with the IP

address of the interface on which the request was received. The agent unicasts the valid packets to the next configured destination. The server responds with a unicast BOOTREPLY addressed to the relay agent closest to the client as indicated by giaddr field. Upon reception of the BOOTREPLY from the server, the agent forwards this reply as broadcast or unicast on the interface form where the BOOTREQUEST was arrived. This interface can be identified by giaddr field.

FASTPATH also supports DHCP relay agent options to identify the source circuit when customers are connected to the Internet with high-speed modem. The relay agent inserts these options when forwarding the request to the server and removes them when sending the reply to the clients.

If an interface has more than one IP address, the relay agent should use the primary IP address configured as its relay agent IP address.

The BOOTP/DHCP Relay Agent folder contains links to the following web pages that configure and display BOOTP/DHCP relay agent:

- BOOTP/DHCP Relay Agent Configuration

4.4.1 BOOTP/DHCP Relay Agent Configuration

Use the BOOTP/DHCP Relay Agent Configuration page to configure and display a BOOTP/DHCP relay agent.

To display the page, click **Routing > BOOTP/DHCP Relay Agent > Configuration** in the navigation tree.

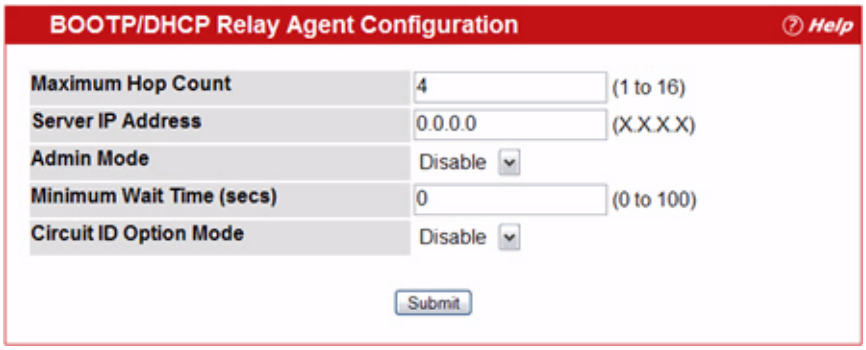

Figure 4-20: BOOTP/DHCP Relay Agent Configuration

Table 4-20: BOOTP/DHCP Relay Agent Configuration Fields

Field	Description
Maximum Hop Count	Enter the maximum number of hops a client request can take before being discarded.

Table 4-20: BOOTP/DHCP Relay Agent Configuration Fields (Continued)

Field	Description
Server IP Address	Enter either the IP address of the BOOTP/DHCP server or the IP address of the next BOOTP/DHCP Relay Agent. ***Note:*** This configuration is deprecated. Use 4.5IP Helper290 to achieve the same functionality.
Admin Mode	Select Enable or Disable from the dropdown menu. When you select Enable, BOOTP/DHCP requests are forwarded to the IP address you entered in the Server IP address field.
Minimum Wait Time (secs)	Enter a time in seconds. This value is compared to the time stamp in the client's request packets, which should represent the time since the client was powered up. Packets are only forwarded when the time stamp exceeds the minimum wait time.
Circuit ID Option Mode	Select Enable or Disable from the dropdown menu. If you select Enable, the relay agent adds Option 82 header packets to the DHCP Request packets before forwarding them to the server, and strips them off while forwarding the responses to the client.

If you make any changes to the page, click **Submit** to apply the changes to the system.

4.5 IP Helper

4.5.1 IP Helper Global Configuration

Use the IP Helper IP Global Configuration page to configure IP Helper globally.

To display the page, click **Routing > IP Helper > Global Configuration** in the navigation tree.

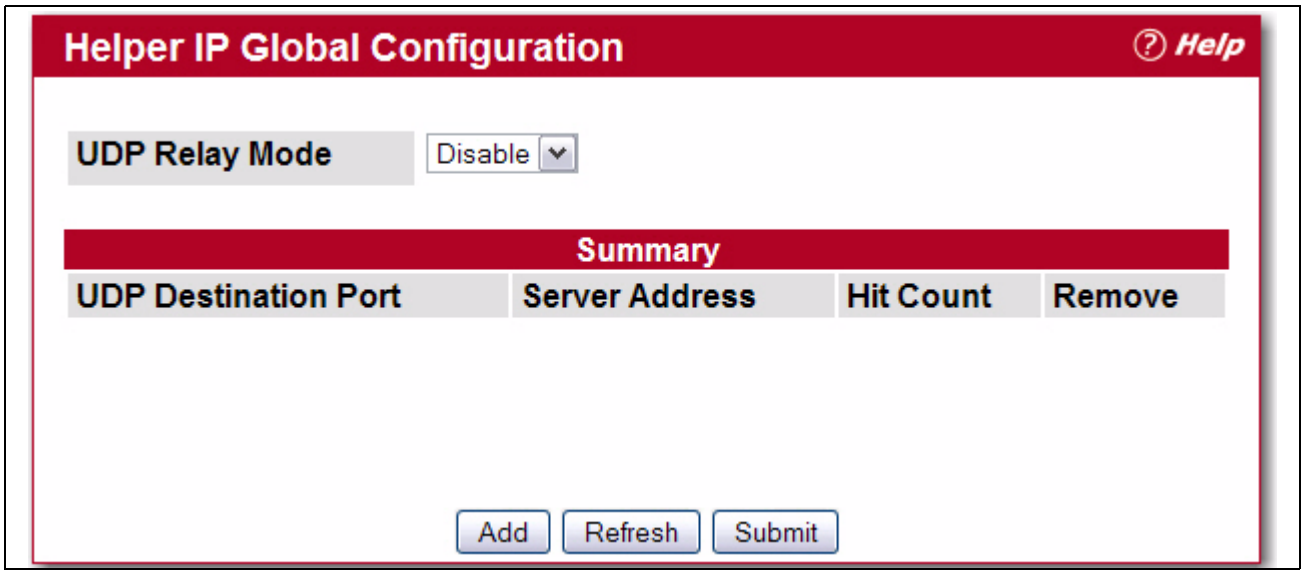

Figure 4-21: IP Helper Global Configuration

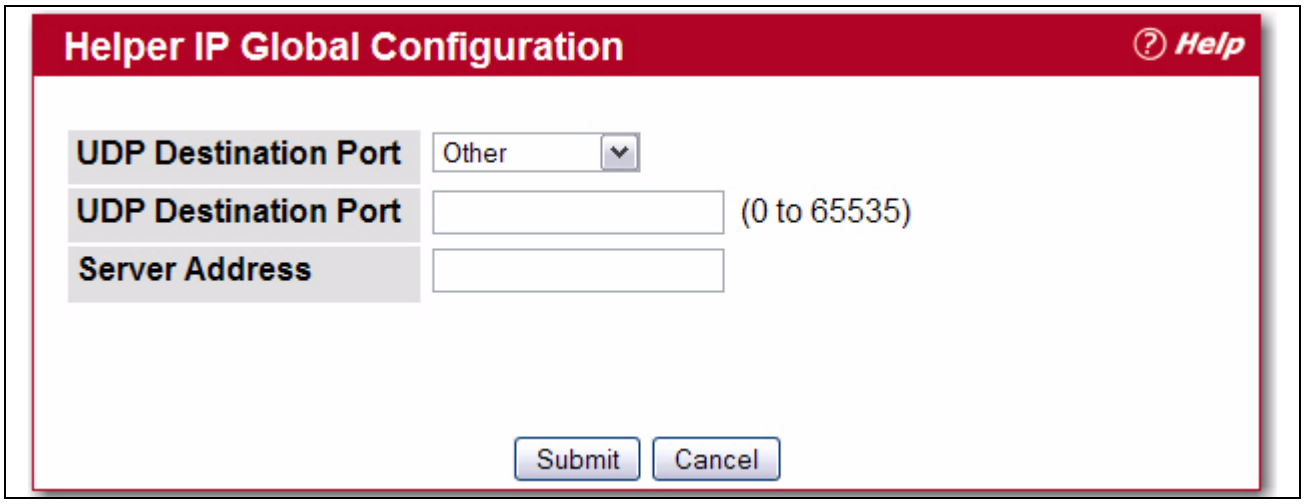

Figure 4-22: IP Helper Global Configuration Add

Table 4-21: IP Helper Global Configuration Add Fields

Field	Description
UDP Destination Port (0-65535)	The destination UDP port ID/Port Name of UDP packets to be relayed. Select the protocol from the menu. If you want to configure other than the listed protocols, select **Other** from the menu. Then user will be prompted with the **UDP Destination Port** field. Select the **DefaultSet** to configure for the relay entry for the default set of protocols.
Server Address	The Server Address to which the packets with the given UDP Destination Port will be relayed.

- Click **Submit** to send the updated configuration to the switch. These changes will not be retained across a power cycle unless a Save configuration is performed.
- Click **Cancel** to cancel the configuration on the screen and reset the data on the screen to the latest value of the switch.

4.5.2 IP Helper Interface Configuration

Use the IP Helper Interface Configuration page to add a IP Helper ID address to the interface.

To display the page, click **Routing > IP Helper > Interface Configuration** in the navigation tree.

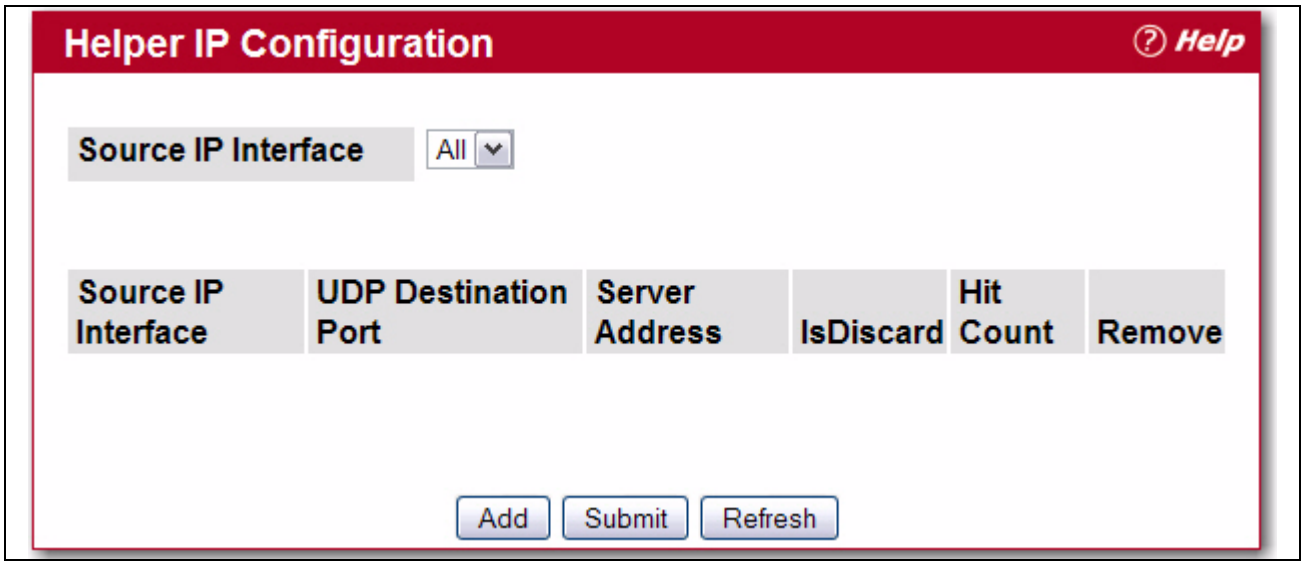

Figure 4-23: IP Helper Global Configuration

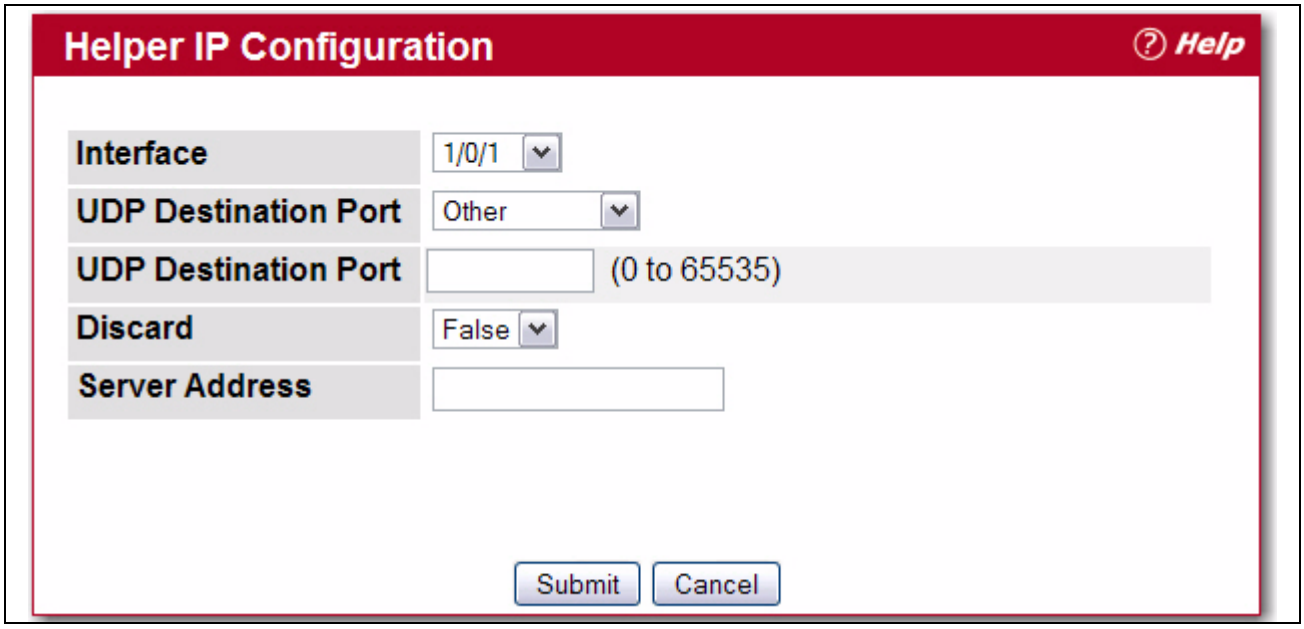

Figure 4-24: IP Helper Global Configuration Add

Table 4-22: IP Helper Global Configuration Add Fields

Field	Description
Source IP Interface	The the interface from the pulldown menu to for which user wants to configure the relay entry.
UDP Destination Port	The the Destination UDP port Name from the pull down menu or configure the port number to configure the Relay Entry on selected interface.
Discard	If set to **True**, packets arriving on the given interface with the given destination UDP port are discarded rather than relayed. Discard entries are used to override global IP helper address entries which otherwise might apply to a packet.
Server Address	The IPv4 address of the server to which packets are relayed for the specific UDP Destination Port.

- Click **Submit** to send the updated configuration to the switch. These changes will not be retained across a power cycle unless a Save configuration is performed.
- Click **Cancel** to cancel the configuration on the screen and reset the data on the screen to the latest value of the switch.

4.5.3 IP Helper Interface Configuration

Use the IP Helper – Helper Statistics page to view IP Helper statistics.

To display the page, click **Routing > IP Helper > Helper Statistics** in the navigation tree.

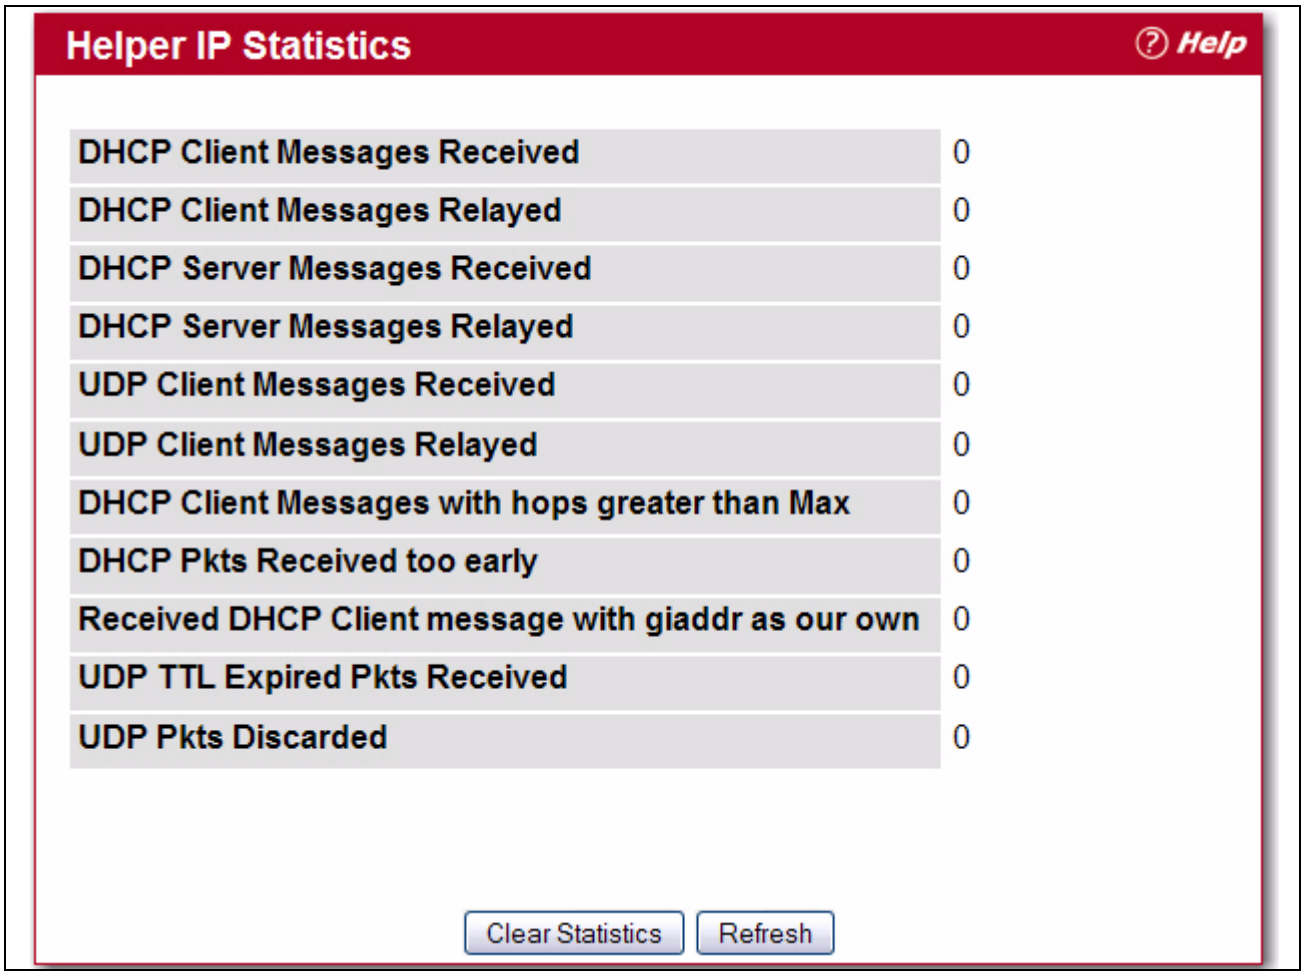

Figure 4-25: IP Helper – Helper Statistics

Table 4-23: IP Helper – Helper Statistics Fields

Field	Description
DHCP Client Messages Received	The number of valid messages received from a DHCP client. The count is only incremented if IP helper is enabled globally, the ingress routing interface is up, and the packet passes a number of validity checks, such as having a TTL >1 and having valid source and destination IP addresses.
DHCP Client Messages Relayed	The number of DHCP client messages relayed to a server. If a message is relayed to multiple servers, the count is incremented once for each server.
DHCP Server Messages Received	The number of DHCP responses received from the DHCP server. This count only includes messages that the DHCP server unicasts to the relay agent for relay to the client.
DHCP Server Messages Relayed	Specifies the number of DHCP server messages relayed to a client.
UDP Client Messages Received	The number of valid UDP packets received. This count includes DHCP messages and all other protocols relayed. Conditions are similar to those for the first statistic in this table.
UDP Client Messages Relayed	The number of UDP packets relayed. This count includes DHCP messages relayed as well as all other protocols. The count is incremented for each server to which a packet is sent.
DHCP Client Messages with hops greater than Max	Specifies the number of DHCP Client Messages with hops greater than Max.
DHCP Pkts Received too early	Specifies the number of DHCP Pkts Received too early.
Received DHCP Client message with giaddr as our own	Specifies the number of DHCP Client messages received with giaddr as our own.
UDP TTL Expired Pkts Received	Specifies the number of UDP packets received with expired TTL.
UDP Pkts Discarded	Specifies the number of UDP packets discarded.

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

4.6 Configuring RIP

RIP is an Interior Gateway Protocol (IGP) based on the Bellman-Ford algorithm and targeted at smaller networks (network diameter no greater than 15 hops). The routing information is propagated in RIP update packets that are sent out both periodically and in the event of a network topology change. On receipt of a RIP update, depending on whether the specified route exists or does not exist in the route table, the router may modify, delete, or add the route to its route table.

The Routing > RIP menu page contains links to the following web pages that configure and display RIP parameters and data:

- RIP Configuration
- RIP Interface Summary
- RIP Interface Configuration
- RIP Route Redistribution Configuration
- RIP Route Redistribution Summary

4.6.1 RIP Configuration

Use the RIP Configuration page to enable and configure or disable RIP in Global mode.

To display the page, click **Routing > RIP > Configuration** in the navigation tree.

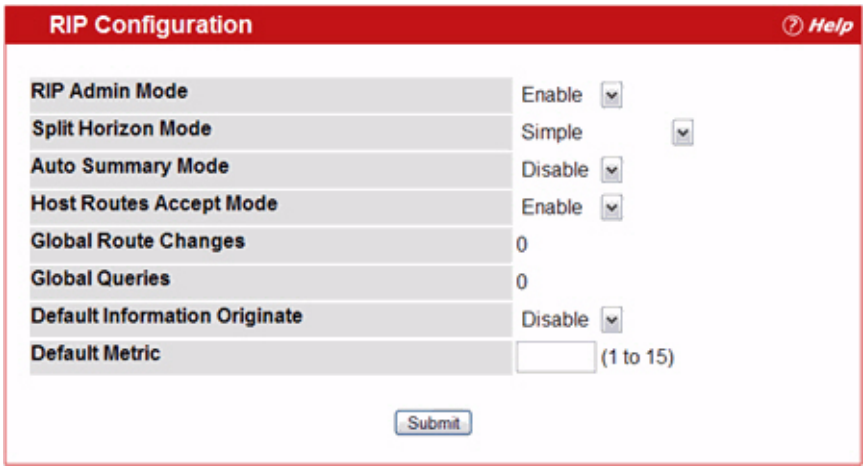

Figure 4-26: RIP Configuration

Table 4-24: RIP Configuration Fields

Field	Description
RIP Admin Mode	Select Enable or Disable from the dropdown menu. If you select Enable, RIP is enabled for the switch. The default is Disable.
Split Horizon Mode	Select None, Simple, or Poison Reverse from the dropdown menu. The default is Simple. Split horizon is a technique for avoiding problems caused by including routes in updates sent to the router from which the route was originally learned. The options are: • **None**: No special processing for this case. • **Simple**: A route is not included in updates sent to the router from which it was learned. • **Poison Reverse**: A route is included in updates sent to the router from which it was learned, but the metric is set to infinity.
Auto Summary Mode	Select Enable or Disable from the dropdown menu. If you select Enable, groups of adjacent routes are summarized into single entries, in order to reduce the total number of entries. The default is Disable.
Host Routes Accept Mode	Select Enable or Disable from the dropdown menu. If you select Enable, the router accepts host routes. The default is Enable.
Global Route Changes	Displays the number of route changes made to the IP Route Database by RIP. This does not include the refresh of a route's age.
Global Queries	Displays the number of responses sent to RIP queries from other systems.
Default Information Originate	When enabled, RIP originates a default route (0.0.0.0/0.0.0.0)
Default Metric	Sets a default for the metric of redistributed routes.This field displays the default metric if one has already been set, or blank if not configured earlier. Valid values are 1 to 15.

If you make changes to the page, click **Submit** to apply the changes to the system.

4.6.2 RIP Interface Summary

Use the RIP Interface Summary page to display RIP configuration status on an interface.

To display the page, click **Routing > RIP > Interface Summary** in the navigation tree.

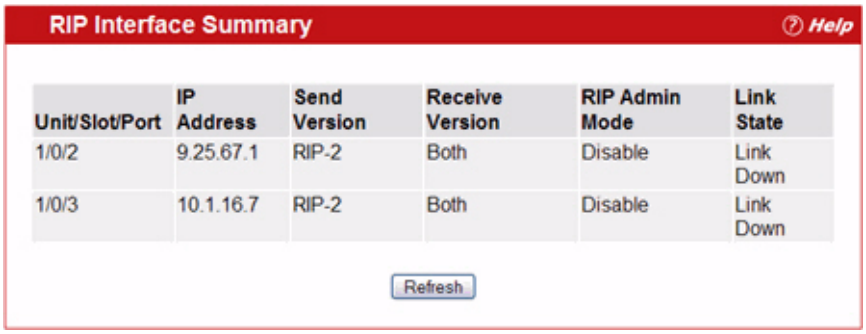

Unit/Slot/Port	IP Address	Send Version	Receive Version	RIP Admin Mode	Link State
1/0/2	9.25.67.1	RIP-2	Both	Disable	Link Down
1/0/3	10.1.16.7	RIP-2	Both	Disable	Link Down

Figure 4-27: RIP Interface Summary

Table 4-25: RIP Interface Summary Fields

Field	Description
Slot/Port	The interface, such as the routing-enabled VLAN on which RIP is enabled.
IP Address	The IP Address of the router interface.
Send Version	Specifies the RIP version to which RIP control packets sent from the interface conform. The default is RIP-2. Possible values are: • **RIP-1**: RIP version 1 packets are sent using broadcast. • **RIP-1c**: RIP version 1 compatibility mode. RIP version 2 formatted packets are transmitted using broadcast. • **RIP-2**: RIP version 2 packets are sent using multicast. • **None**: RIP control packets are not transmitted.
Receive Version	Specifies which RIP version control packets are accepted by the interface. The default is Both. Possible values are: • **RIP-1**: only RIP version 1 formatted packets are received. • **RIP-2**: only RIP version 2 formatted packets are received. • **Both**: packets are received in either format. • **None**: no RIP control packets are received.
RIP Admin Mode	Specifies whether RIP is Enabled or Disabled on the interface.
Link State	Specifies whether the RIP interface is up or down.

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

4.6.3 RIP Interface Configuration

Use the RIP Interface Configuration page to enable and configure or to disable RIP on a specific interface.

To display the page, click **Routing > RIP > Interface Configuration** in the navigation tree.

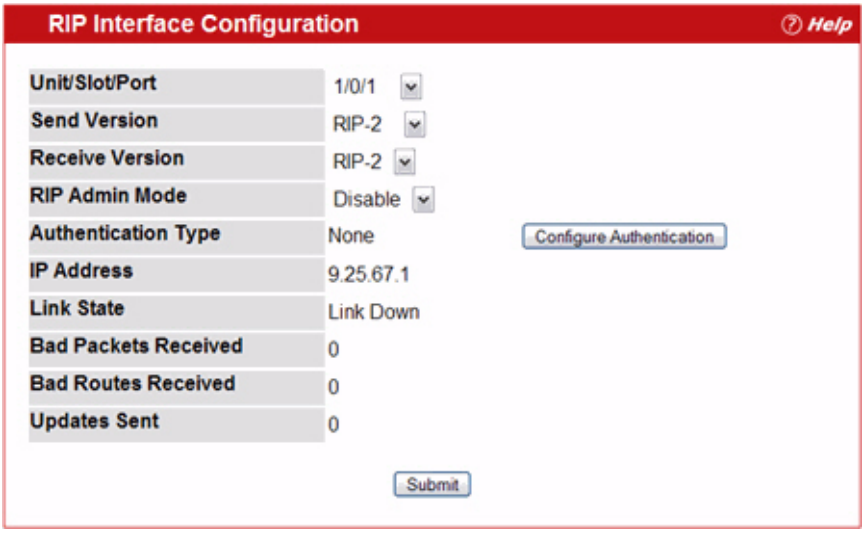

Figure 4-28: RIP Interface Configuration

Table 4-26: RIP Interface Configuration Fields

Field	Description
Slot/Port	Select the interface for which data is to be configured from the dropdown menu.
Send Version	RIP Version that router sends with its routing updates. The default is RIP-2. Possible values are: • **RIP-1**: send RIP version 1 formatted packets via broadcast. • **RIP-1c**: RIP version 1 compatibility mode. Send RIP version 2 formatted packets via broadcast. • **RIP-2**: send RIP version 2 packets using multicast. • **None**: no RIP control packets are sent.
Receive Version	RIP Version of the routing updates that the router must accept. The default is Both. Possible values are: • **RIP-1**: accept only RIP version 1 formatted packets. • **RIP-2**: accept only RIP version 2 formatted packets. • **Both**: accept packets in either format. • **None**: no RIP control packets is accepted.
RIP Admin Mode	Select Enable or Disable from the dropdown menu. Before you enable RIP version 1 or version 1c on an interface, you must first enable network directed broadcast mode on the corresponding interface. The default value is Disable.
Authentication Type	You may select an authentication type other than None by clicking the **Modify** button. You then see a new screen, where you can select the authentication type from the dropdown menu. Possible values are: • **None**: This is the initial interface state. If you select this option from the dropdown menu on the second screen you are returned to the first screen without any authentication protocols being run. • **Simple**: If you select Simple you are prompted to enter an authentication key. This key is included, in the clear, in the RIP header of all packets sent on the network. All routers on the network must be configured with the same key. • **Encrypt**: If you select Encrypt you are prompted to enter both an authentication key and an authentication ID. Encryption uses the MD5 Message-Digest algorithm. All routers on the network must be configured with the same key and ID.
IP Address	Displays the IP Address of the router interface.

Table 4-26: RIP Interface Configuration Fields (Continued)

Field	Description
Link State	Specifies whether the RIP interface is up or down.
Bad Packets Received	Displays the number of RIP packets that were found to be invalid or corrupt.
Bad Routes Received	Displays the number of routes, in valid RIP packets, which were ignored for any reason, e.g., the number of triggered RIP updates actually sent on this interface. This explicitly does NOT include full updates sent containing new information.
Updates Sent	Displays the number of route updates sent.

4.6.3.1 Configuring the RIP Interface

1. Open the **RIP Interface Configuration** page.
2. Specify the interface for which data is to be configured.
3. Enter data into the fields as needed.
4. To change the **Authentication Type**, click **Configure Authentication** to configure different Authentication Types.

   The page refreshes and displays the OSPF Interface Authentication Configuration page.

Figure 4-29: RIP Interface Authentication Configuration

5. Select the type of authentication to use.

   If you select Simple or Encrypt as the authentication, the screen refreshes, and additional fields display. Enter the required information into the new fields.
6. Click **Submit** to apply the changes to the system and return to the **RIP Interface Configuration** page.
7. To cancel the authentication configuration and return to the **RIP Interface Configuration** page, click **Cancel**.

4.6.4 RIP Route Redistribution Configuration

Use the RIP Route Redistribution Configuration page to configure which routes are redistributed to other routers using RIP. The allowable values for each fields are displayed next to the field. If any invalid values are entered, an alert message is displayed with the list of all the valid values.

To display the page, click **Routing > RIP > Route Redistribution Configuration** in the navigation menu.

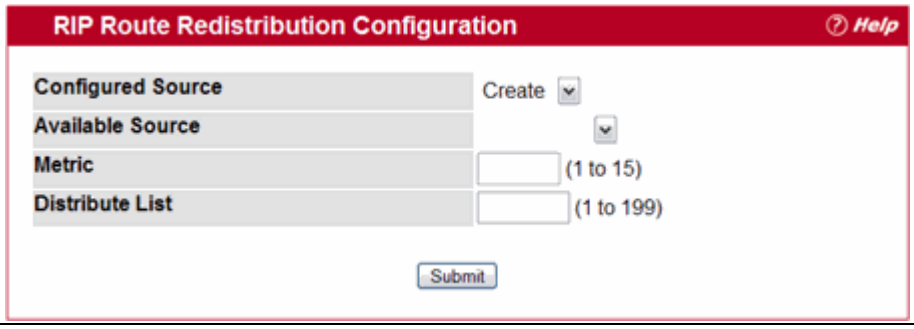

Figure 4-30: RIP Route Redistribution Configuration

Table 4-27: RIP Route Redistribution Configuration Fields

Field	Description
Configured Source	If any Source Routes have already been configured for redistribution by RIP, they appear in the dropdown menu. Otherwise, the only available option is Create, which allows you to configure an available source route. Select an existing route to view or modify its parameters.
Available Source	The dropdown menu is populated by only those Source Routes that have not previously been configured for redistribution by RIP. This field is available only if you select Create as the Configured Source. Possible values are: • Static: The route was manually configured. • Connected: The route was determined automatically because the host is directly connected. • RIP: The route was determined through RIP. • BGP: The route was determined through BGP.
Metric	Sets the metric value to be used as the metric of redistributed routes. This field displays the metric if the source was pre-configured and can be modified. The valid values are 1 to 15.
Match	Displays only when OSPF is selected as the Available Source type. Use these fields to further specify the types of OSPF-learned routes that are subject to redistribution using RIP: • Internal routes • External Type 1 Routes • External Type 2 Routes • NSSA External Type 1 Routes • NSSA External Type 2 Routes
Distribute List	Enter the ACL ID for an access list that filters the routes to be redistributed by the destination protocol. Only permitted routes are redistributed.

You configure ACLs through the pages under **QoS > Access Control Lists** > **IP Access Control Lists**. When used for route filtering, the only fields in an access list that get used are:

- Source IP Address and netmask
- Destination IP Address and netmask
- Action (Permit or Deny)

All other fields (source and destination port, precedence, tos, etc.) are ignored.

The source IP address is compared to the destination IP address of the route. The source IP netmask in the access list rule is treated as a wildcard mask, indicating which bits in the source IP address must match the

destination address of the route. (Note that a 1 in the mask indicates a Don't Care in the corresponding address bit.)

When an access list rule includes a destination IP address and netmask (an extended access list), the destination IP address is compared to the network mask of the destination of the route. The destination netmask in the access list serves as a wildcard mask, indicating which bits in the route's destination mask are significant for the filtering operation.

- If you make changes to the page, click **Submit** to apply the changes to the system.
- To delete a configured route, click **Delete**.

4.6.5 RIP Route Redistribution Summary

Use the RIP Route Redistribution Summary page to display Route Redistribution configurations.

To display the page, click **Routing > RIP > Route Redistribution Summary** in the navigation menu.

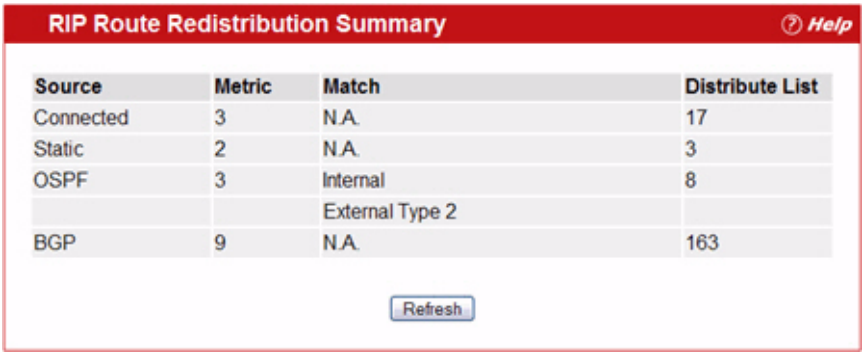

Figure 4-31: RIP Route Redistribution Summary

Table 4-28: RIP Route Redistribution Summary Fields

Field	Description
Source	The protocol used to obtain the route.
Metric	The Metric of redistributed routes for the given source route. Displays Unconfigured when not configured.
Match	List of Routes redistributed when OSPF is selected as Source. If OSPF is not the source, the field is N.A. or "not applicable." The list may include one or more of: • Internal • External 1 • External 2 • NSSA-External 1 • NSSA-External 2
Distribute List	The Access List that filters the routes to be redistributed by the Destination Protocol. If the Distribute List is not configured, the field is blank.

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

4.7 Router Discovery

The Router Discovery protocol is used by hosts to identify operational routers on the subnet. Router Discovery messages are of two types: “Router Advertisements” and “Router Solicitations.” The protocol mandates that every router periodically advertise the IP Addresses it is associated with. Hosts listen for these advertisements and discover the IP Addresses of neighboring routers.

The **Routing > Router Discovery** folder contains links to the following web pages that configure and display Router Discovery data:

- Router Discovery Configuration
- Router Discovery Status

4.7.1 Router Discovery Configuration

Use the Router Discovery Configuration page to enter or change Router Discovery parameters.

To display the page, click **Routing > Router Discovery > Configuration** in the navigation tree.

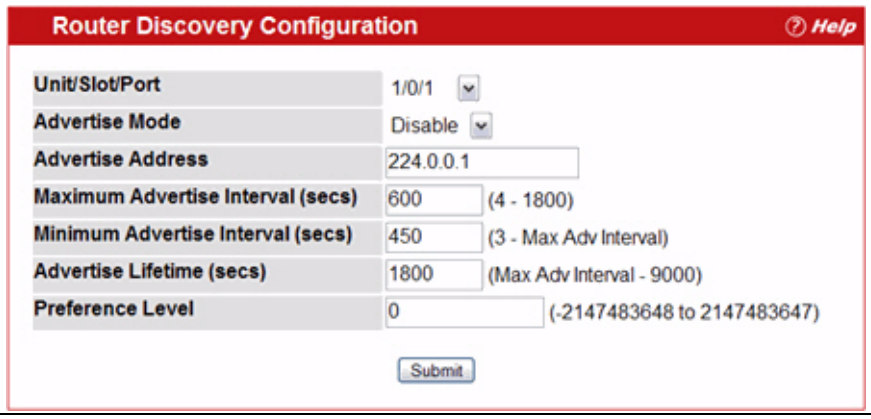

Figure 4-32: Router Discovery Configuration

Table 4-29: Router Discovery Configuration Fields

Field	Description
Slot/Port	Select the router interface for which data is to be configured.
Advertise Mode	Select Enable or Disable from the dropdown menu. If you select Enable, Router Advertisements are transmitted from the selected interface.
Advertise Address	Enter the IP Address to be used to advertise the router.
Maximum Advertise Interval (secs)	Enter the maximum time (in seconds) allowed between router advertisements sent from the interface.
Minimum Advertise Interval (secs)	Enter the minimum time (in seconds) allowed between router advertisements sent from the interface.
Advertise Lifetime (secs)	Enter the value (in seconds) to be used as the lifetime field in router advertisements sent from the interface. This is the maximum length of time that the advertised addresses are to be considered as valid router addresses by hosts.
Preference Level	Specify the preference level of the router as a default router relative to other routers on the same subnet. Higher numbered addresses are preferred. You must enter an integer.

If you make any changes to the page, click **Submit** to apply the changes to the system.

4.7.2 Router Discovery Status

Use the Router Discovery Status page to display Router Discovery data for each port.

To display the page, click **Routing > Router Discovery > Status** in the navigation tree.

Router Discovery Status Help

Unit/Slot/Port	Advertise Mode	Advertise Address	Maximum Advertise Interval (secs)	Minimum Advertise Interval (secs)	Advertise Lifetime (secs)	Preference Level
1/0/2	Enable	224.0.0.1	600	450	1800	0
1/0/3	Disable	224.0.0.1	600	450	1800	0
1/0/6	Disable	224.0.0.1	600	450	1800	0
1/0/7	Disable	224.0.0.1	600	450	1800	0
1/0/8	Enable	255.255.255.255	600	450	1800	0
1/0/11	Disable	224.0.0.1	600	450	1800	0
1/0/12	Disable	224.0.0.1	600	450	1800	0
1/0/13	Disable	224.0.0.1	600	450	1800	0
1/0/14	Disable	224.0.0.1	600	450	1800	0
1/0/16	Disable	224.0.0.1	600	450	1800	0
1/0/17	Disable	224.0.0.1	600	450	1800	0
1/0/18	Disable	224.0.0.1	600	450	1800	0
1/0/19	Disable	224.0.0.1	600	450	1800	0
1/0/20	Disable	224.0.0.1	600	450	1800	0
1/0/21	Disable	224.0.0.1	600	450	1800	0
1/0/22	Disable	224.0.0.1	600	450	1800	0
1/0/23	Disable	224.0.0.1	600	450	1800	0
1/0/24	Disable	224.0.0.1	600	450	1800	0
1/0/25	Disable	224.0.0.1	600	450	1800	0
1/0/26	Disable	224.0.0.1	600	450	1800	0
1/0/27	Disable	224.0.0.1	600	450	1800	0
1/0/28	Disable	224.0.0.1	600	450	1800	0

Refresh

Figure 4-33: Router Discovery Status

Table 4-30: Router Discovery Status Fields

Field	Description
Slot/Port	The router interface for which data is displayed.
Advertise Mode	The values are Enable or Disable. Enable denotes that Router Discovery is enabled on that interface.
Advertise Address	The IP Address used to advertise the router.

Table 4-30: Router Discovery Status Fields (Continued)

Field	Description
Maximum Advertise Interval(secs)	The maximum time (in seconds) allowed between router advertisements sent from the interface.
Minimum Advertise Interval(secs)	The minimum time (in seconds) allowed between router advertisements sent from the interface.
Advertise Lifetime(secs)	The value (in seconds) used as the lifetime field in router advertisements sent from the interface. This is the maximum length of time that the advertised addresses are to be considered as valid router addresses by hosts.
Preference Level	The preference level of the router as a default router relative to other routers on the same subnet. Higher numbered addresses are preferred.

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

4.8 Router

The Routing > Router folder contains links to the following web pages that configure and display route tables:

- Route Table
- Best Routes Table
- Configured (Static) Routes
- Route Preferences Configuration

4.8.1 Route Table

The route table manager collects routes from multiple sources: static routes, RIP routes, OSPF routes, BGP routes, 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 (see 4.8.2Best Routes Table305 for more information).

To display the page, click **Routing > Router > Route Table** in the navigation tree.

Figure 4-34: Route Table

Table 4-31: Route Table Fields

Field	Description
Total Number of Routes	The total number of routes in the route table.
Network Address	The IP route prefix for the destination.
Subnet Mask	Also referred to as the subnet/network mask, this indicates the portion of the IP interface address that identifies the attached network.
Protocol	This field tells which protocol created the specified route. The possibilities are one of the following: • Local • Static • Default • OSPF Intra • OSPF Inter • OSPF Type-1 • OSPF Type-2 • RIP
Next Hop Slot/Port	The outgoing router interface to use when forwarding traffic to the destination.
Next Hop IP Address	The outgoing router IP address to use when forwarding traffic to the next router (if any) in the path towards the destination. The next router is always one of the adjacent neighbors or the IP address of the local interface for a directly attached network.

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

4.8.2 Best Routes Table

The route table manager collects routes from multiple sources: static routes, RIP routes, OSPF routes, BGP routes, local routes. The route table manager may learn multiple routes to the same destination from multiple sources. In that case, the route table manager selects the route with the lowest route preference value to use for forwarding to that destination. Use the Best Routes Table page to display the best routes from the routing table. To view all routes, including multiple routes to the same destination, see 4.8.1Route Table 304.

To display the page, click **Routing > Router > Best Routes Table** in the navigation tree.

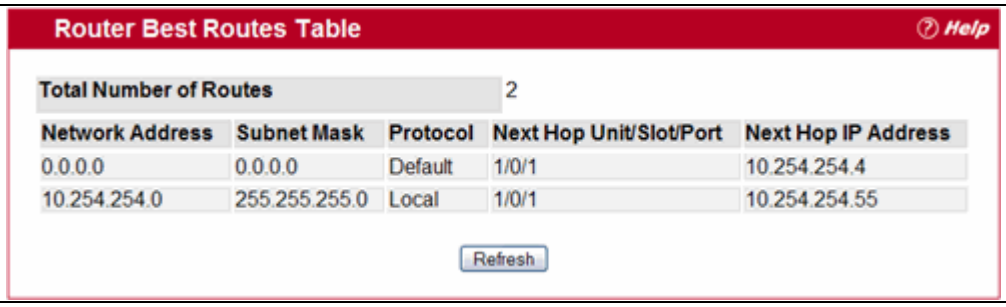

Figure 4-35: Best Routes Table

Table 4-32: Best Routes Table Fields

Field	Description
Total Number of Routes	The total number of routes in the route table.
Network Address	The IP route prefix for the destination.
Subnet Mask	Also referred to as the subnet/network mask, this indicates the portion of the IP interface address that identifies the attached network.
Protocol	This field tells which protocol created the specified route. The possibilities are one of the following: • Local • Static • Default • OSPF Intra • OSPF Inter • OSPF Type-1 • OSPF Type-2 • RIP
Next Hop Slot/Port	The outgoing router interface to use when forwarding traffic to the destination.
Next Hop IP Address	The outgoing router IP address to use when forwarding traffic to the next router (if any) in the path towards the destination. The next router is always one of the adjacent neighbors or the IP address of the local interface for a directly attached network.

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

4.8.3 Configured (Static) Routes

Use the Configured Routes page to create and display static routes.

To display the page, click **Routing > Router > Configured Routes** in the navigation tree.

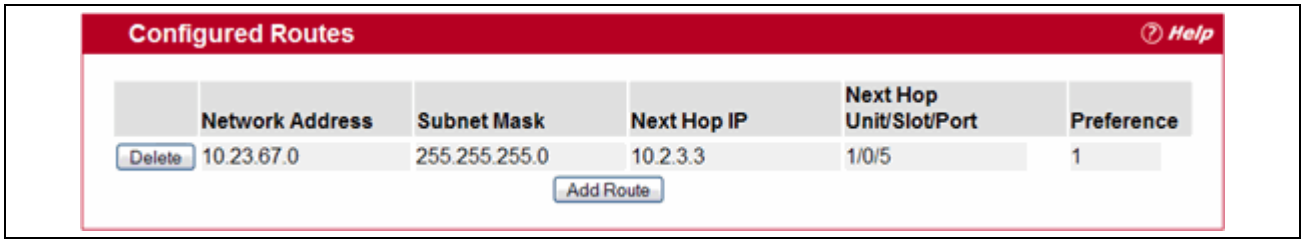

Figure 4-36: Configured Routes

Table 4-33: Configured Routes Fields

Field	Description
Network Address	The IP route prefix for the destination.
Subnet Mask	Also referred to as the subnet/network mask, this indicates the portion of the IP interface address that identifies the attached network.
Next Hop IP	The next hop router address to use when forwarding traffic to the destination.
Next Hop Unit/Slot Port	The outgoing interface to use when forwarding traffic to the destination. For static reject routes it would be Null0.
Preference	The preferences configured for the added routes.

4.8.3.1 Adding a Static Route

1. Open the Configured Routes page.
2. Click **Add**.

   The **Router Route Entry Configuration** page displays:

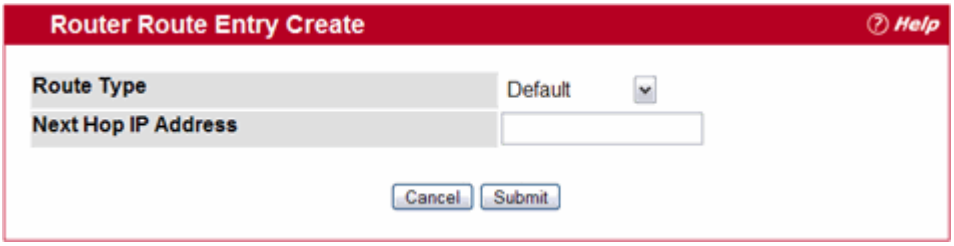

Figure 4-37: Create Default Route Entry

3. Next to **Route Type**, select **Default** route, **Static** or **Static Reject** from the menu.

   Default: Enter the default gateway address in the **Next Hop IP Address** field.

   Static: Enter values for **Network Address**, **Subnet Mask**, **Next Hop IP Address**, and **Preference**.

   Static Reject: Packets to these destinations will be dropped.

   If you select Static as the route type, the screen refreshes and additional fields appear, as Figure 4-39 shows.

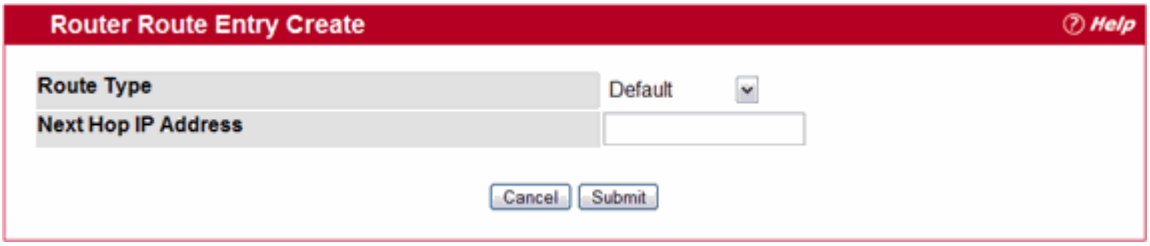

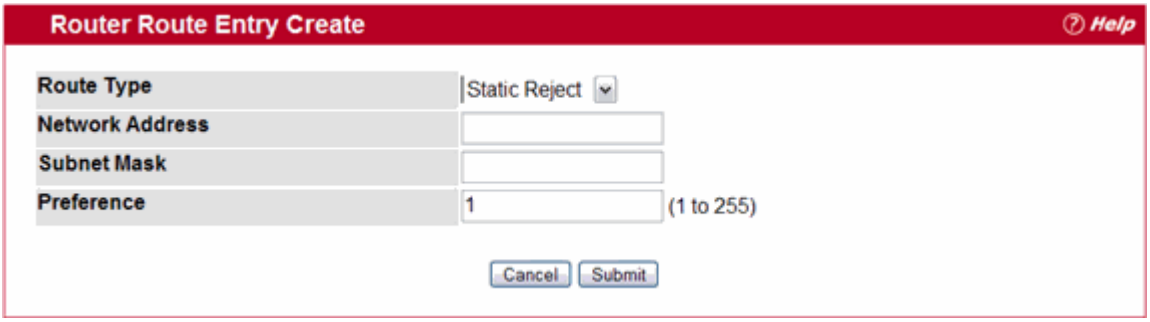

Figure 4-38: Create Static Route Entry

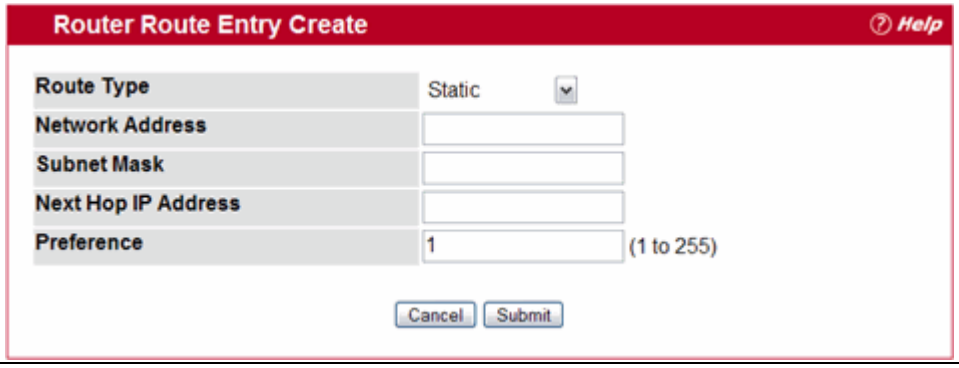

Figure 4-39: Create Static Route Entry

Table 4-34: Route Entry Create Fields

Field	Description
Network Address	Specify the IP route prefix for the destination from the dropdown menu. In order to create a route, a valid routing interface must exist and the next hop IP Address must be on the same network as the routing interface. Routing interfaces are created on the **IP Interface Configuration** page. Valid next hop IP Addresses can be viewed on the **Route Table** page.
Subnet Mask	Also referred to as the subnet/network mask, this indicates the portion of the IP interface address that identifies the attached network.
Protocol	This field tells which protocol created the specified route. Possible values are: • Local • Static • Default • OSPF Intra • OSPF Inter • OSPF Type-1 • OSPF Type-2 • RIP
Next Hop Slot/Port	The outgoing router interface to use when forwarding traffic to the destination.

Table 4-34: Route Entry Create Fields (Continued)

Field	Description
Next Hop IP Address	The outgoing router IP address to use when forwarding traffic to the next router (if any) in the path towards the destination. The next router is always one of the adjacent neighbors or the IP address of the local interface for a directly attached network. When creating a route, the next hop IP must be on the same network as the routing interface. Valid next hop IP Addresses can be seen on the 'Route Table' page.
Metric	Administrative cost of the path to the destination. If no value is entered, default is 1. The range is 0 - 255.This field is present only when creating a static route.
Preference	Specifies a preference value for the configured next hop.
Route Type	Specifies whether the route is to be a Default route or a Static route.

4. Click **Submit**.

   The new route is added, and you are returned to the Configured Routes page.

4.8.3.2 Deleting a Route

Click **Delete** to remove a configured route.

4.8.4 Route Preferences Configuration

Use the Route Preferences Configuration page to configure the default preference for each protocol. These values are arbitrary values that range from 1 to 255, and are independent of route metrics. Most routing protocols use a route metric to determine the shortest path known to the protocol, independent of any other protocol. Routes with a preference of 255 are not used for forwarding.

The best route to a destination is chosen by selecting the route with the lowest preference value. When there are multiple routes to a destination, the preference values are used to determine the preferred route.

To display the page, click **Routing > Router > Route Preferences Configuration** in the navigation tree.

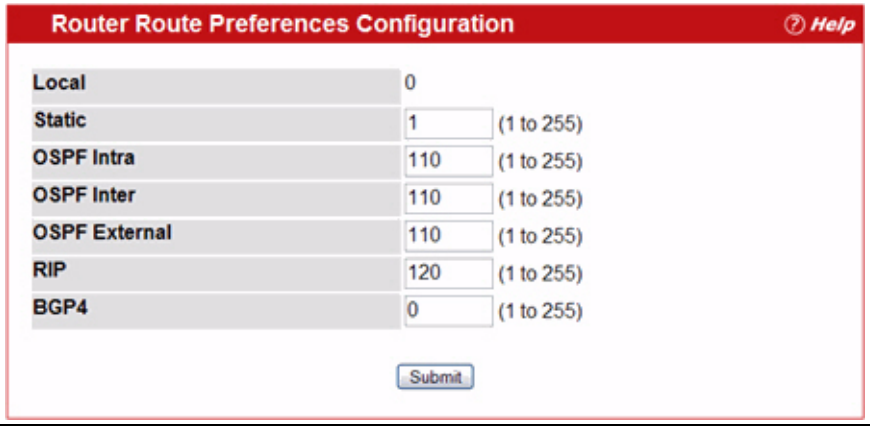

Figure 4-40: Route Preferences Configuration

Table 4-35: Route Preferences Configuration Fields

Field	Description
Local	This field displays the local route preference value of 0. This value is not configurable.
Static	The static route preference value in the router. The default value is 1. The range is 1 to 255.
OSPF Intra	The OSPF intra route preference value in the router. The default value is 8. The range is 1 to 255. The OSPF specification (RFC 2328) requires that preferences must be given to the routes learned via OSPF in the following order: intra < inter < type-1 < type-2.
OSPF Inter	The OSPF inter route preference value in the router. The default value is 10. The range is 1 to 255. The OSPF specification (RFC 2328) requires that preferences must be given to the routes learned via OSPF in the following order: intra < inter < type-1 < type-2.
OSPF Type-1	The OSPF type-1 route preference value in the router. The default value is 13. The range is 1 to 255. The OSPF specification (RFC 2328) requires that preferences must be given to the routes learned via OSPF in the following order: intra < inter < type-1 < type-2.
OSPF Type-2	The OSPF type-2 route preference value in the router. The default value is 150. The range is 1 to 255. The OSPF specification (RFC 2328) requires that preferences must be given to the routes learned via OSPF in the following order: intra < inter < type-1 < type-2.
RIP	The RIP route preference value in the router. The default value is 15. The range is 1 to 255.
BGP4	The BGP4 route preference. The default value is 0. The range is 1 to 255.

If you make changes to the page, click **Submit** to apply the changes to the system.

4.9 VLAN Routing

You can configure FASTPATH software with some ports supporting VLANs and some supporting routing. You can also configure the software to allow traffic on a VLAN to be treated as if the VLAN were a router port.

When a port is enabled for bridging (default) rather than routing, all normal bridge processing is performed for an inbound packet, which is then associated with a VLAN. Its MAC Destination Address (MAC DA) and VLAN ID are used to search the MAC address table. If routing is enabled for the VLAN, and the MAC DA of an inbound unicast packet is that of the internal bridge-router interface, the packet is routed. An inbound multicast packet is forwarded to all ports in the VLAN, plus the internal bridge-router interface, if it was received on a routed VLAN.

Since a port can be configured to belong to more than one VLAN, VLAN routing might be enabled for all of the VLANs on the port, or for a subset. VLAN Routing can be used to allow more than one physical port to reside on the same subnet. It could also be used when a VLAN spans multiple physical networks, or when additional segmentation or security is required. This section shows how to configure FASTPATH software to support VLAN routing. A port can be either a VLAN port or a router port, but not both. However, a VLAN port may be part of a VLAN that is itself a router port.

The Routing > VLAN Routing menu page contains links to the following web pages that allow you to configure and display VLAN Routing parameters and data:

- VLAN Routing Configuration

- VLAN Routing Summary

4.9.1 VLAN Routing Configuration

Use the VLAN Routing Configuration page to configure VLAN Routing interfaces on the system.

To display the page, click **Routing > Router > VLAN Routing Configuration** in the navigation tree.

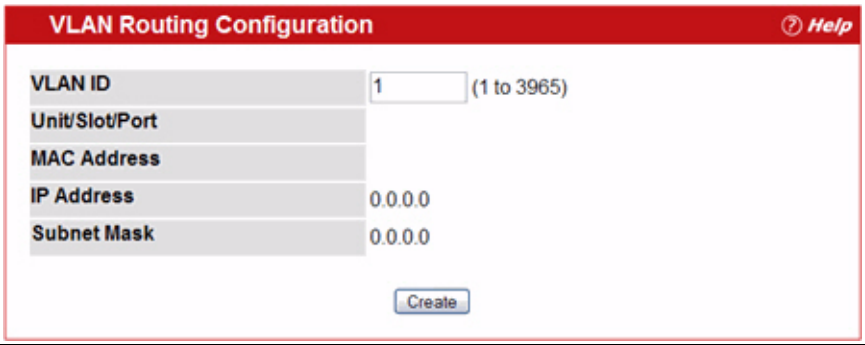

Figure 4-41: VLAN Routing Configuration

Table 4-36: VLAN Routing Configuration Fields

Field	Description
VLAN ID	Enter the ID of a VLAN to configure for VLAN Routing. Initially, the field will display the ID of the first VLAN. After you enter a new VLAN ID and click **Create**, the non-configurable data will be displayed.
Slot/Port	The logical slot and port number assigned to the VLAN Routing Interface.
MAC Address	The MAC Address assigned to the VLAN Routing Interface.
IP Address	The configured IP Address of the VLAN Routing Interface. Note that if a VLAN is created and the IP address is not configured, the page by default shows an IP address of 0.0.0.0. To configure the IP address, go to **IP > Interface Configuration**.
Subnet Mask	The configured Subnet Mask of the VLAN Routing Interface. This is 0.0.0.0 when the VLAN Routing Interface is first configured and must be entered on the IP Interface Configuration page.

4.9.1.1 Creating a VLAN Routing Interface

1. Enter a new VLAN ID in the **VLAN ID** field.
2. Click Create.

   The page refreshes and displays the interface and MAC address assigned to the new VLAN. The interface is in Slot/Port notation. The IP address and Subnet Mask fields are 0.0.0.0.

Note...

Be sure to note the interface Unit/Slot assignment so that you select the correct interface to configure from the Interface Configuration page.

3. in the navigation menu, click **Routing > IP > Interface Configuration**.

4. Select the interface assigned to the VLAN.

   The IP address and Subnet Mask fields are 0.0.0.0 by default.
5. Enter the IP address and subnet mask for the VLAN, and configure any other interface settings.
6. Click **Submit** to apply the settings to the VLAN routing interface.
7. Navigate to the **Routing > VLAN Routing > Summary** page to view the new VLAN in the table.

4.9.1.2 Deleting a VLAN Router Interface

Click **Delete** to delete the selected VLAN routing interface.

4.9.2 VLAN Routing Summary

Use the VLAN Routing Summary page to display information about the VLAN Routing interfaces configured on the system.

To display the page, click **Routing > Router > VLAN Routing Summary** in the navigation tree.

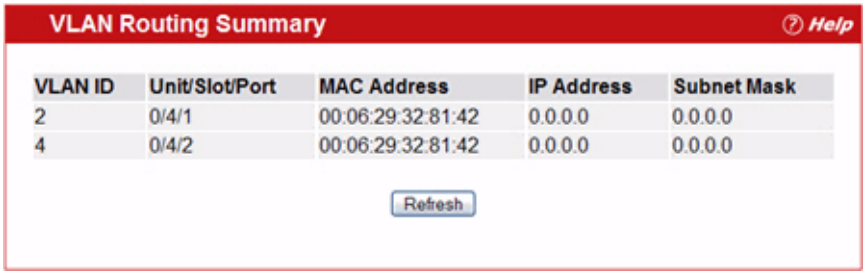

Figure 4-42: VLAN Routing Summary

Table 4-37: VLAN Routing Summary Fields

Field	Description
VLAN ID	The ID of the VLAN whose data is displayed in the current table row.
Slot/Port	The logical slot and port number assigned to the VLAN Routing Interface.
MAC Address	The MAC Address assigned to the VLAN Routing Interface.
IP Address	The configured IP Address of the VLAN Routing Interface. Note that if a VLAN is created and the IP address is not configured, the page by default shows an IP address of 0.0.0.0. To configure the IP address, go to **IP→ Interface Configuration**.
Subnet Mask	The configured Subnet Mask of the VLAN Routing Interface. This is 0.0.0.0 when the VLAN Routing Interface is first configured and must be entered on the IP Interface Configuration page.

4.10 Virtual Router Redundancy Protocol (VRRP)

The Virtual Router Redundancy protocol is designed to handle default router failures by providing a scheme to dynamically elect a backup router. The driving force was to minimize “black hole” periods due to the failure of the default gateway router during which all traffic directed towards it is lost until the failure is

detected. Though static configuration of default routes is popular, such an approach is susceptible to a single point of failure when the default router fails. VRRP advocates the concept of a “virtual router” associated with one or more IP Addresses that serve as default gateways. In the event that the VRRP Router controlling these IP Addresses (formally known as the Master) fails, the group of IP Addresses and the default forwarding role is taken over by a Backup VRRP Router.

The **Routing > VRRP** folder contains links to the following web pages that configure and display VRRP parameters and data:

- VRRP Configuration
- Virtual Router Configuration
- Virtual Router Status
- Virtual Router Statistics

4.10.1 VRRP Configuration

Use the VRRP Configuration page to enable or disable the administrative status of a virtual router.

To display the page, click **Routing > VRRP > Configuration** in the navigation tree.

Figure 4-43: VRRP Configuration

Table 4-38: VRRP Configuration

Field	Description
Admin Mode	This sets the administrative status of VRRP in the router to active or inactive. Select Enable or Disable from the dropdown menu. The default is Disable.

If you change the administrative mode, click **Submit** to apply the changes to the system.

4.10.2 Virtual Router Configuration

Use the Virtual Router Configuration page to create a new virtual router or to configure an existing one.

To display the page, click **Routing > VRRP > Virtual Router Configuration** in the navigation tree.

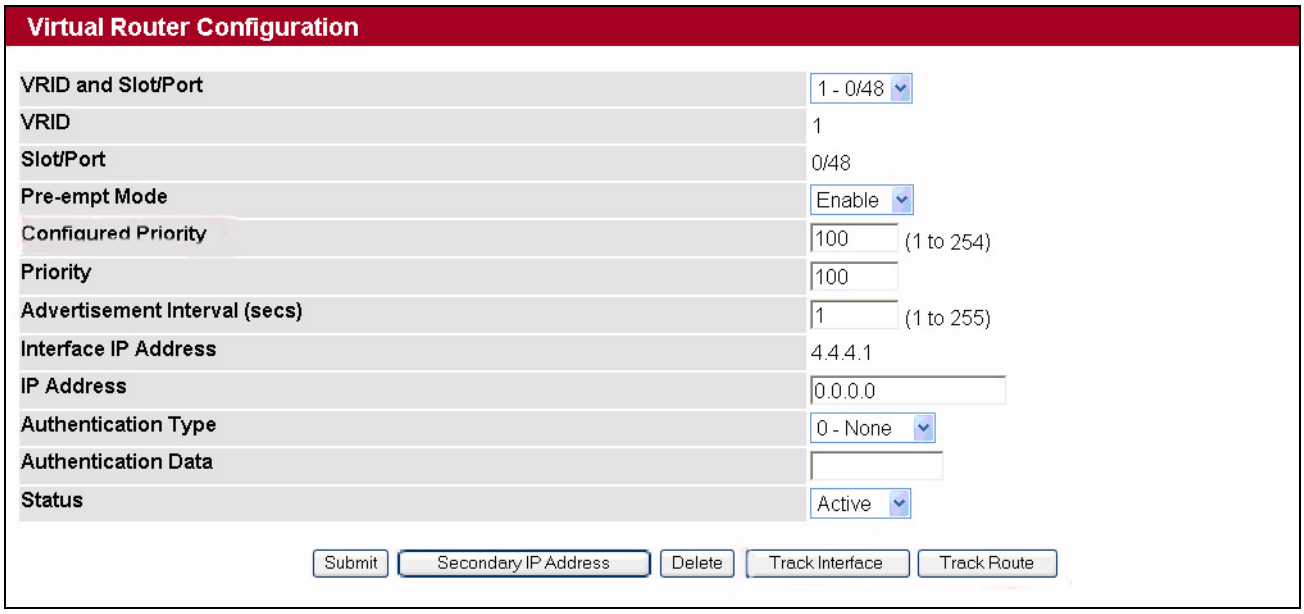

Figure 4-44: Virtual Router Configuration

Table 4-39: Virtual Router Configuration Fields

Field	Description
VRID and Slot/Port	Select **Create** from the dropdown menu to configure a new Virtual Router, or select one of the existing Virtual Routers, listed by interface number and VRID.
VRID	This field is only configurable if you are creating new Virtual Router, in which case enter the VRID in the range 1 to 255.
Slot/Port	This field is only configurable if you are creating new Virtual Router, in which case select the interface for the new Virtual Router from the dropdown menu.
Pre-empt Mode	Select Enable or Disable from the dropdown menu. If you select Enable, a backup router preempts the master router if it has a priority greater than the master virtual router's priority, provided that the master is not the owner of the virtual router's IP address. The default is Enable.
Configured Priority	Enter the priority value to be used by the VRRP router in the election for the master virtual router. If the Virtual IP Address is the same as the interface IP Address, the priority gets set to 255 no matter what you enter. If you enter a priority of 255 when the Virtual and interface IP Addresses are not the same, the priority gets set to the default value of 100.
Priority	The operational priority of the VRRP router. This is relative to the configured priority. The operational priority depends upon the configured priority, and the priority decrements configured through the tracking process.
Advertisement Interval (secs)	Enter the time, in seconds, between the transmission of advertisement packets by this virtual router. Enter a number between 1 and 255. The default value is 1 second.
Interface IP Address	Indicates the IP Address associated with the selected interface.

Table 4-39: Virtual Router Configuration Fields (Continued)

Field	Description
IP Address	Enter the IP Address associated with the Virtual Router. The default is 0.0.0.0, which you must change prior to clicking **Create**.
Authentication Type	Select the type of Authentication for the Virtual Router from the dropdown menu. The default is None. The choices are: • **0-None**: No authentication is performed. • **1-Simple**: Authentication is performed using a text password.
Authentication Data	If you selected simple authentication, enter the password.
Status	Select active or inactive from the dropdown menu to start or stop the operation of the Virtual Router. The default is inactive.

- Click **Submit** to apply the new configuration and cause the change to take effect. These changes will not be retained across a power cycle unless a Save configuration is performed.
- Click **Secondary IP Address** to proceed to the Secondary IP Address configuration page.
- Click **Delete** to delete the selected Virtual Router. Note that the router cannot be deleted if there are secondary addresses configured.
- Click **Track Interface** to proceed to the VRRP Track Interface configuration page.
- Click **Track Route** to proceed to the VRRP Track Route configuration page.

4.10.2.1 Configuring a Secondary VRRP Address

To configure a secondary VRRP address, first configure one IP address (the primary address) for the VR. Then, you can add multiple secondary addresses to that interface.

- Click **Submit** to apply the new configuration and cause the change to take effect. These changes will not be retained across a power cycle unless a Save configuration is performed.
- Click **Delete** to delete the selected secondary IP address.
- Click **Cancel** to return to the Virtual Router Configuration page.

4.10.2.2 Creating a New Virtual Router

1. From the **Virtual Router Configuration** page, select **Create** from the VRID and Slot/Port dropdown menu.
2. Specify the VRID, the virtual router address, and the interface for the new virtual router.
3. Define the remaining fields as needed.
4. Click **Create** to apply the changes to the system.

   The new virtual router is saved, and the device is updated.

4.10.2.3 Modifying a Virtual Router

To modify the settings for an existing virtual router, select its ID from the VRID and Slot/Port dropdown menu and change the fields as needed. Click **Submit** to apply the changes to the system.

4.10.2.4 VRRP Interface Tracking Configuration

Use VRRP Interface Tracking to track a specific interface IP state within the router that can alter the priority level of a virtual router for a VRRP group. An exception to this is, if that VRRP group is the IP address owner,

its priority is fixed at 255 and cannot be reduced through the tracking process.

To display the page, click **Routing > VRRP > Virtual Router Configuration** in the navigation tree, then click the **Track Interface** button.

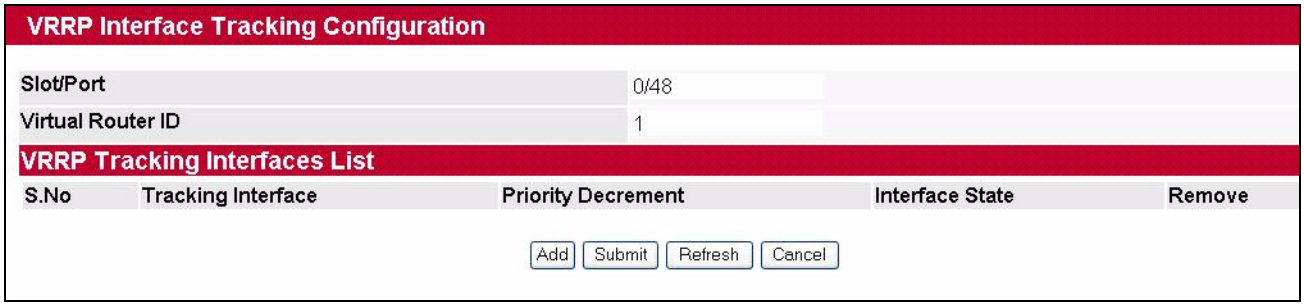

Figure 4-45: VRRP Interface Tracking Configuration

Table 4-40: VRRP Interface Tracking Configuration Fields

Field	Description
Slot/Port	The interface associated with the Virtual Router ID.
Virtual Router ID	The Virtual Router ID for which data is to be displayed.
S. No	The serial number for this row.
Tracking Interface	The Tracked Interface for which data is to be displayed.
Priority Decrement	The priority decrement for the tracked interface. The valid range is 1 to 254. The default value is 10.
Interface State	The IP state of the tracked interface.
Remove	Removes the selected Tracking Interface from the VRRP tracked list.

- Click **Add** to proceed to the VRRP Interface Tracking page.
- Click **Submit** to apply the new configuration. Configuration changes take effect immediately. These changes will not be retained across a power cycle unless a Save configuration is performed.
- Click **Refresh** to refresh the page with the most current data from the switch.
- Click **Cancel** to return to the Virtual Router Configuration page.

4.10.2.5 VRRP Interface Tracking

Use the VRRP Interface Tracking page to add an interface to the tracking list.

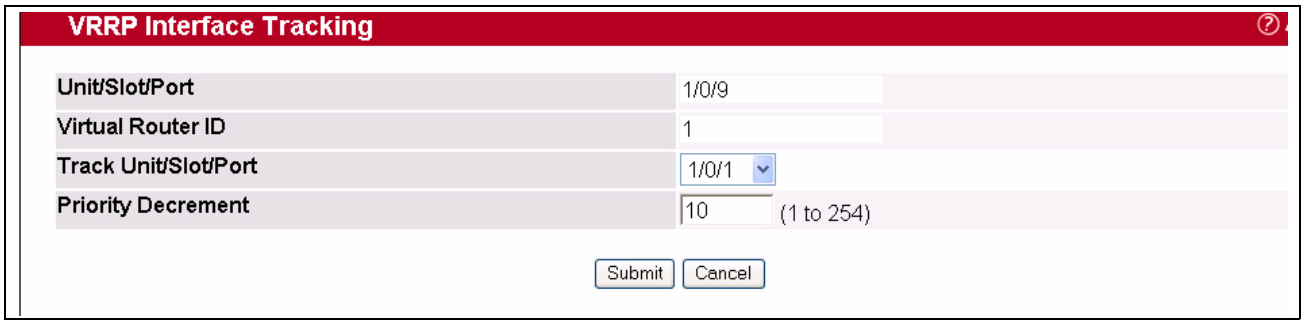

Figure 4-46: VRRP Interface Tracking

Table 4-41: VRRP Track Interface Fields

Field	Description
Slot/Port	The interface associated with the Virtual Router ID.
Virtual Router ID	The Virtual Router ID for which data is to be displayed.
Track Slot/Port	Displays all routing interfaces which are not yet tracked for this Virtual Router ID and interface configuration. Exceptions to this: loopback and tunnels could not be tracked.
Priority Decrement	The priority decrement for the tracked interface. The valid range is 1-254. The default value is 10.

- Click **Submit** to send the updated configuration to the switch. Configuration changes take effect immediately. These changes will not be retained across a power cycle unless a Save configuration is performed.
- Click **Cancel** to return to the VRRP Interface Tracking Configuration page.

4.10.2.6 VRRP Route Tracking Configuration

Use VRRP Route Tracking Configuration to track specific route IP states within the router that can alter the priority level of a virtual router for a VRRP group.

To display the page, click **Routing > VRRP > Virtual Router Configuration** in the navigation tree, then click the **Track Route** button.

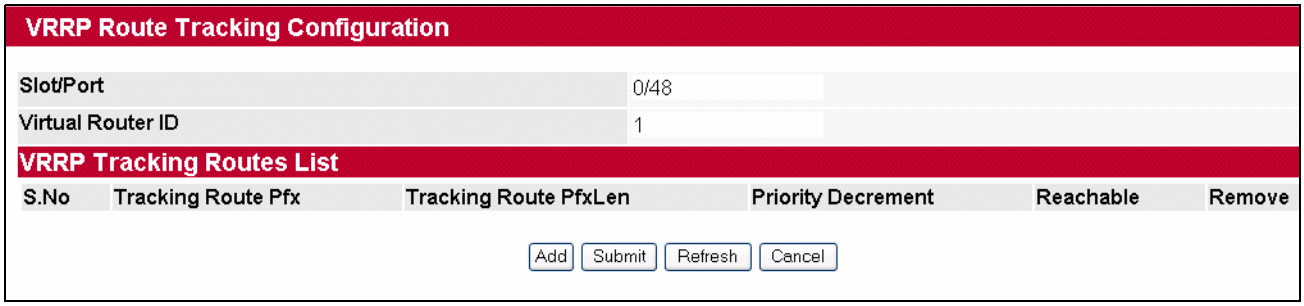

Figure 4-47: VRRP Route Tracking Configuration

Table 4-42: VRRP Route Tracking Configuration Fields

Field	Description
Slot/Port	The interface associated with the Virtual Router ID.
Virtual Router ID	The Virtual Router ID for which tracking data is to be displayed.
S. No	The serial number for this row.
Tracking Route Pfx	The prefix of the tracked route.

Table 4-42: VRRP Route Tracking Configuration Fields (Continued)

Field	Description
Tracking Route PfxLen	The prefix length of the tracked route.
Priority Decrement	Enter the priority decrement for the tracked route. The valid range is 1-254. The default value is 10.
Reachable	The reachability of the tracked route.
Remove	Removes the selected tracking routes from the VRRP tracked list.

- Click **Add** to proceed to the VRRP Route Tracking page.
- Click **Submit** to send the updated configuration to the switch. Configuration changes take effect immediately. These changes will not be retained across a power cycle unless a Save configuration is performed.
- Click **Refresh** to refresh the page with the most current data from the switch.
- Click **Cancel** to return to the Virtual Router Configuration page.

4.10.2.7 VRRP Route Tracking

Use the VRRP Route Tracking page to add a route into the tracking list.

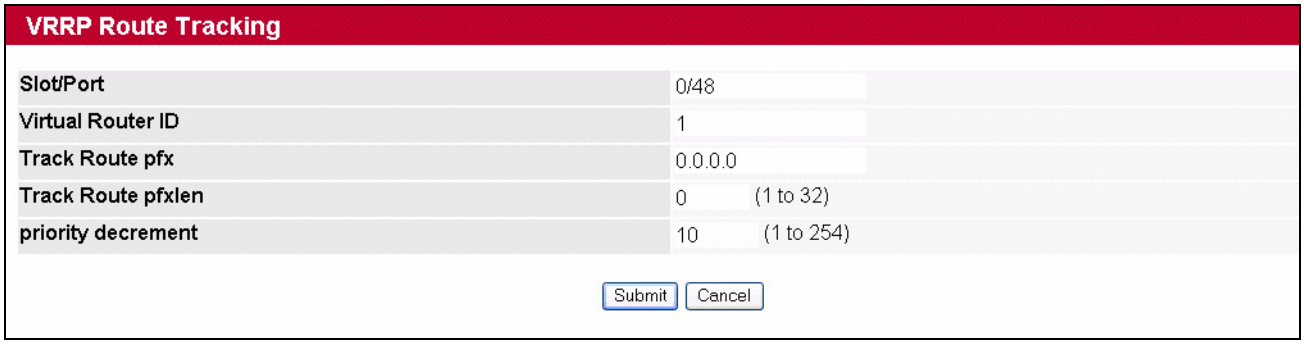

Figure 4-48: VRRP Route Tracking

Table 4-43: VRRP Route Tracking Fields

Field	Description
Slot/Port	The Interface associated with the Virtual Router ID.
Virtual Router ID	The Virtual Router ID for which data is to be displayed.
Track Route Pfx	The prefix of the route.
Track Route PfxLen	The prefix length of the route.
Priority Decrement	The priority decrement for the route. The valid range is 1-254. The default value is 10.

- Click **Submit** to send the updated configuration to the switch. Configuration changes take effect immediately. These changes will not be retained across a power cycle unless a Save is performed.
- Click **Cancel** to return to the VRRP Route Tracking Configuration page.

4.10.3 Virtual Router Status

Use the Virtual Router Status page to display virtual router status.

To display the page, click **Routing > VRRP > Virtual Router Status** in the navigation tree.

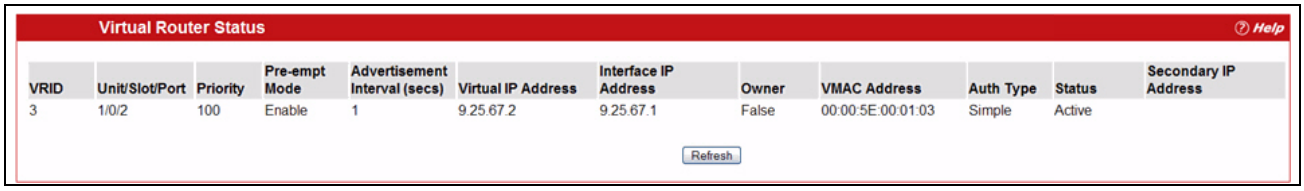

Figure 4-49: Virtual Router Status

Table 4-44: Virtual Router Status Fields

Field	Description
VRID	Virtual Router Identifier.
Slot/Port	Indicates the interface associate with the VRID.
Priority	The priority value used by the VRRP router in the election for the master virtual router.
Pre-empt Mode	• **Enable**: If the Virtual Router is a backup router, it preempts the master router if it has a priority greater than the master virtual router's priority provided that the master is not the owner of the virtual router IP address. • **Disable**: If the Virtual Router is a backup router it does not preempt the master router even if its priority is greater.
Advertisement Interval(secs)	The time, in seconds, between the transmission of advertisement packets by this virtual router.
Virtual IP Address	The IP Address associated with the Virtual Router.
Interface IP Address	The actual IP Address associated with the interface used by the Virtual Router.
Owner	Set to True if the Virtual IP Address and the Interface IP Address are the same, otherwise set to False. If this parameter is set to True, the Virtual Router is the owner of the Virtual IP Address, and always wins an election for master router when it is active.
VMAC Address	The virtual MAC Address associated with the Virtual Router, composed of a 24-bit organizationally unique identifier, the 16-bit constant identifying the VRRP address block and the 8-bit VRID. The Virtual MAC address is: 00:00:5e:00:01:XX, where XX is the VRID.
Auth Type	The type of authentication in use for the Virtual Router • None: Specifies that the authentication type is none. • Simple: Specifies that the authentication type is a simple text password.
State	The current state of the Virtual Router: • Initialize • Master • Backup
Status	The current status of the Virtual Router: • Inactive • Active
Secondary IP Address	A secondary VRRP address configured for the primary VRRP.

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

4.10.4 Virtual Router Statistics

Use the Virtual Router Statistics page to display statistics for a specified virtual router.

To display the page, click **Routing > VRRP > Virtual Router Statistics** in the navigation tree.

Figure 4-50 shows the fields on the **Virtual Router Statistics** page for a switch that has one or more virtual routers configured.

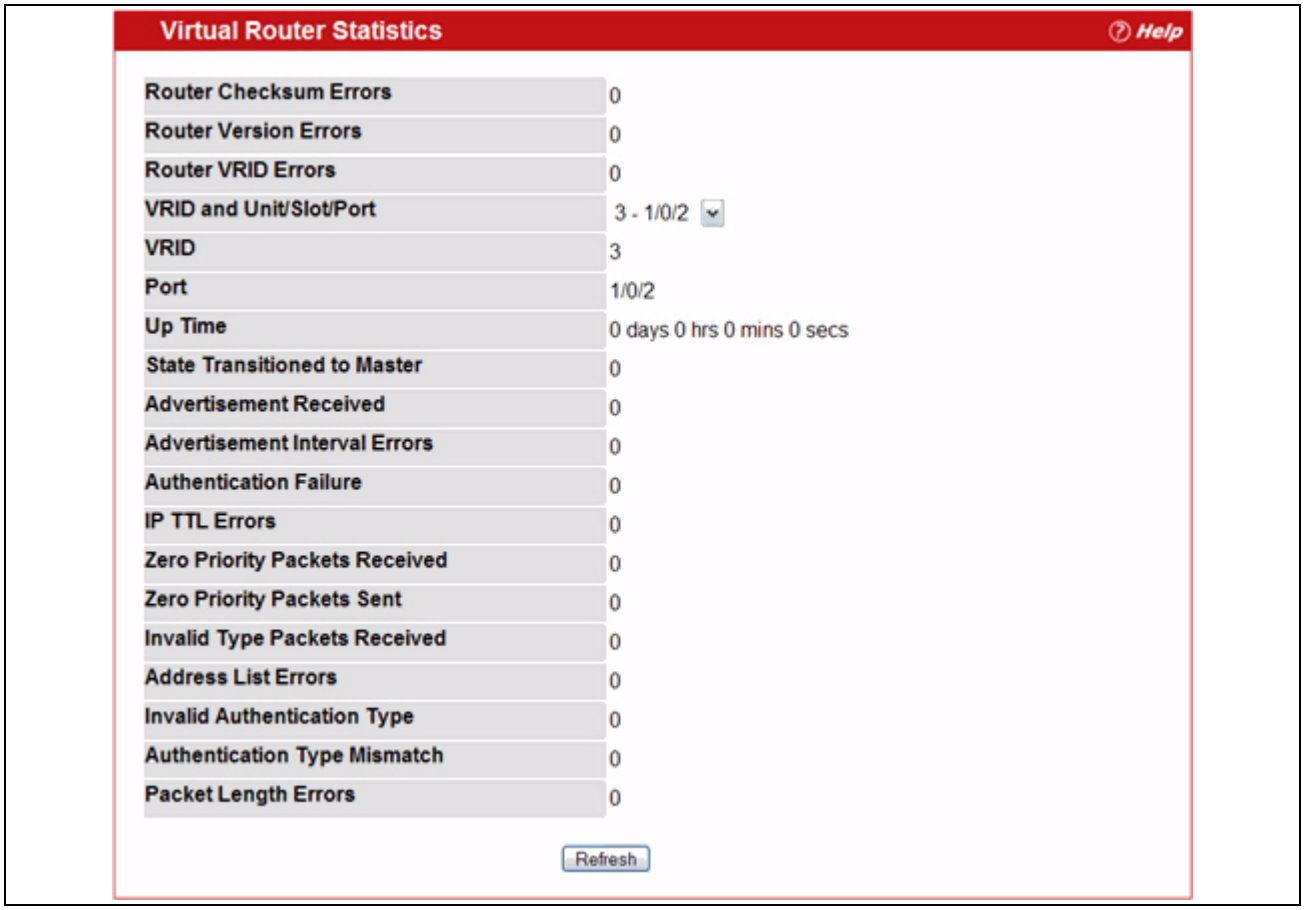

Figure 4-50: Virtual Router Statistics—Virtual Router Configured

The Virtual Router Statistics page contains the fields listed below. Many of the fields display only when there is a valid VRRP configuration.

Table 4-45: Virtual Router Statistics Fields

Field	Description
Router Checksum Errors	The total number of VRRP packets received with an invalid VRRP checksum value.
Router Version Errors	The total number of VRRP packets received with an unknown or unsupported version number.
Router VRID Errors	The total number of VRRP packets received with an invalid VRID for this virtual router.
VRID and Slot/Port	Select the existing Virtual Router, listed by interface number and VRID, for which you want to display statistical information.
VRID	the VRID for the selected Virtual Router.
Slot/Port	The interface for the selected Virtual Router.
Up Time	The time, in days, hours, minutes and seconds, that has elapsed since the virtual router transitioned to the initialized state.
State Transitioned to Master	The total number of times that this virtual router's state has transitioned to Master.
Advertisement Received	The total number of VRRP advertisements received by this virtual router.
Advertisement Interval Errors	The total number of VRRP advertisement packets received for which the advertisement interval was different than the one configured for the local virtual router.
Authentication Failure	The total number of VRRP packets received that did not pass the authentication check.
IP TTL Errors	The total number of VRRP packets received by the virtual router with IP TTL (Time-To-Live) not equal to 255.
Zero Priority Packets Received	The total number of VRRP packets received by the virtual router with a priority of 0.
Zero Priority Packets Sent	The total number of VRRP packets sent by the virtual router with a priority of 0.
Invalid Type Packets Received	The number of VRRP packets received by the virtual router with an invalid value in the Type field.
Address List Errors	The total number of packets received for which the address list does not match the locally configured list for the virtual router.
Invalid Authentication Type	The total number of packets received with an unknown authentication type.
Authentication Type Mismatch	The total number of packets received with an authentication type different to the locally configured authentication method.
Packet Length Errors	The total number of packets received with a packet length less than the length of the VRRP header.

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

4.11 Tunnels

FASTPATH software provides for the creation, deletion, and management of tunnel interfaces. These are dynamic interfaces that are created and deleted via user-configuration.

FASTPATH supports configured IPv6 over IPv4 tunnels to facilitate the transition of IPv4 networks to IPv6 networks. With configured tunnels, the user specifies the endpoints of the tunnel. Tunnels operate as point-to-point links.

Loopback Interfaces (p. 285). Delete the following. I think the reference to the network port is just confusing. You may need to rework the start of the following sentence.

The Routing > Tunnels folder contains links to the following web pages that configure and display tunnel parameters and data:

- Tunnels Configuration
- Tunnel Summary

4.11.1 Tunnels Configuration

Use the Tunnels Configuration page to create, configure, or delete a tunnel.

To display the page, click **Routing > Tunnels > Configuration** in the navigation tree.

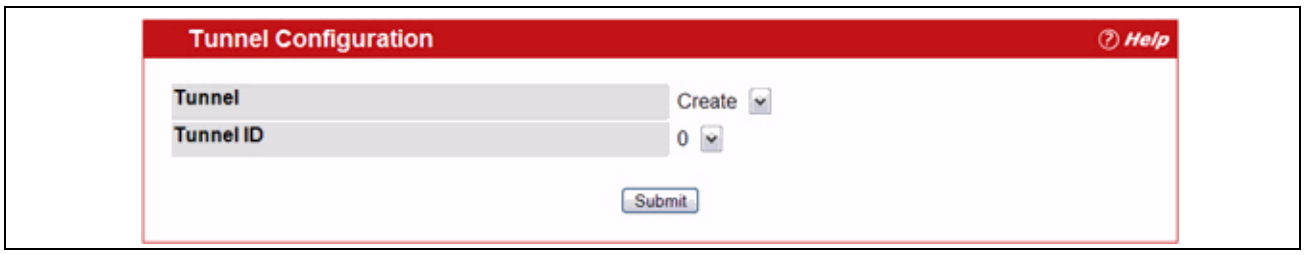

Figure 4-51: Tunnels Configuration

If any tunnels are configured on the system, or if you select Create and click **Submit**, the screen displays with additional fields, as shown in Figure 4-52.

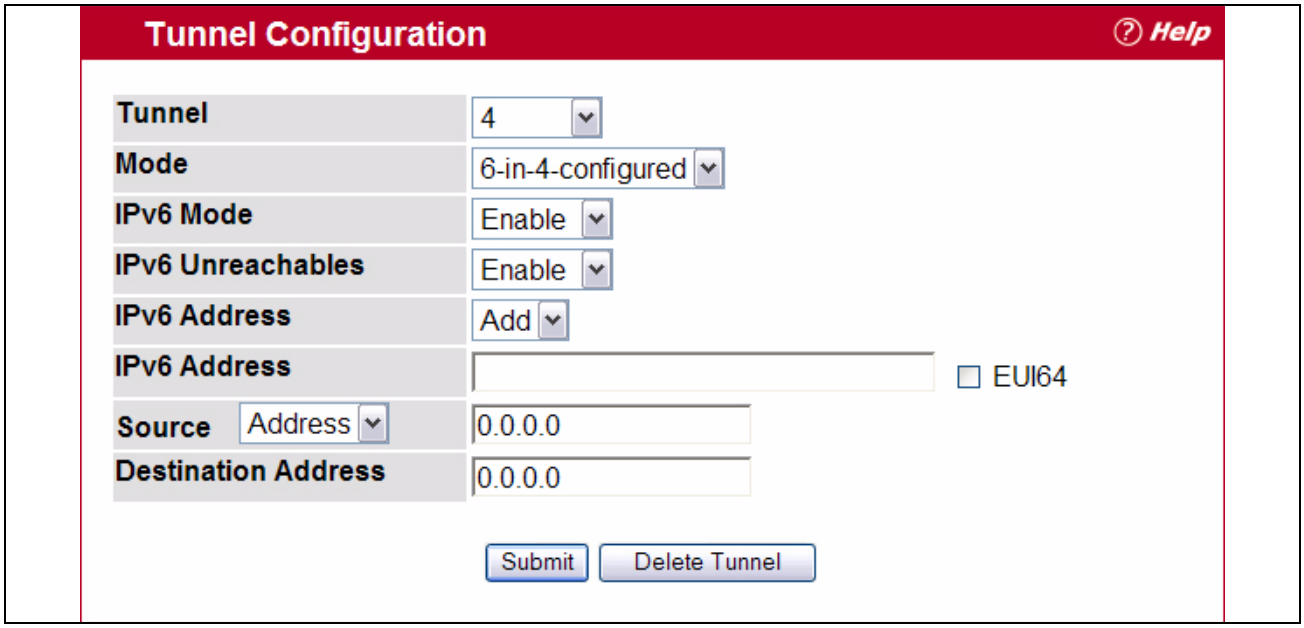

Figure 4-52: Tunnels Configuration—Create Tunnel

Table 4-46: Tunnels Configuration Fields

Field	Description
Tunnel	Use the dropdown menu to select from the list of currently configured tunnel IDs. Create is also a valid choice if the maximum number of tunnel interfaces has not been created.
Tunnel ID	When Create is chosen from the tunnel selector this list of available tunnel IDs becomes visible. You must select a tunnel ID to associate with the new tunnel and click **Apply Changes** before the remaining fields on the page display.
Mode	Selector for the Tunnel mode. The supported modes are 6-in-4-configured and 6-to-4.
IPv6 Mode	Enable IPv6 on this interface using the IPv6 address. This option is only configurable prior to specifying an explicit IPv6 address. When IPv6 mode is enabled, the interface is capable of IPv6 operation without a global address. In this case, an eui-64 based link-local address is used. Select **Enable** or **Disable**. The default value is **Disable**.
IPv6 Unreachables	Specify the Mode of sending ICMPv6 Destination Unreachables on this interface. If disabled, then this interface will not send ICMPv6 Destination Unreachables. By default, the IPv6 Destination Unreachables mode is **Enable**.
IPv6 Address	Select an IPv6 addresses for the selected Tunnel interface. Add is also a valid choice if the maximum number of addresses has not been configured.
IPv6 Address	When Add is chosen from the IPv6 Address selector this IPv6 address input field becomes visible. Address must be entered in the format prefix/length.
Source	Select the desired source, IPv4 Address or Interface. If Address is selected, the source address for this tunnel must be entered in dotted decimal notation. If Interface is selected the source interface for this tunnel must be selected. The address associated with the selected interface is used as the source address.
Destination Address	The IPv4 destination address for this tunnel in dotted decimal notation.

You also have the option to specify the 64-bit extended unique identifier (EUI-64).

4.11.1.1 Creating a New Tunnel

1. Open the **Tunnels Configuration** page.
2. Select **Create** from the **Tunnel** dropdown menu.
3. Specify an ID to use in the **Tunnel ID** field.
4. Click **Submit**.
5. Configure the fields as needed.
6. Enter desired values in the remaining fields.
7. Click **Submit**.

   The new tunnel is saved, and the device is updated.

4.11.1.2 Modifying an Existing Tunnel

1. Open the **Tunnels Configuration** page.
2. Specify the tunnel to modify in the **Tunnel** dropdown menu.
3. Change field values as desired in the remaining fields.
4. Click **Apply Changes**.

   The new configuration is saved, and the device is updated.

4.11.1.3 Removing a Tunnel

1. Open the **Tunnels Configuration** page.
2. Specify the tunnel to remove in the **Tunnel** dropdown menu.
3. Click **Delete Tunnel**.

   The tunnel is deleted, and the device is updated.

4.11.2 Tunnel Summary

Use the Tunnel Summary page to display a summary of configured tunnels.

To display the page, click **Routing > Tunnels > Summary** in the navigation tree.

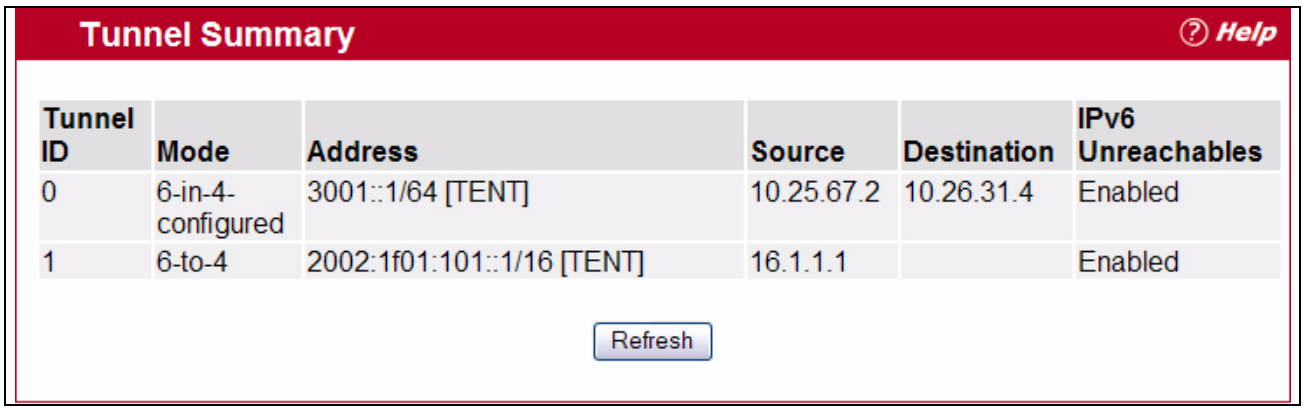

Figure 4-53: Tunnel Summary

Table 4-47: Tunnel Summary Fields

Field	Description
Tunnel ID	The Tunnel ID.
Mode	The corresponding mode of the Tunnel.
Address	The IPv6 Address(es) of the Tunnel.
Source	The corresponding Tunnel Source Address. In the case where an interface has been configured both the interface and the address are displayed. If the source interface has no address configured the text 'unconfigured' is displayed in place of the address.
Destination	The corresponding Tunnel Destination Address.
IPv6 Unreachables	Indicates whether the sending of ICMPv6 Destination Unreachable messages is enabled or disabled. If disables, then this interface will not send ICMPv6 Destination Unreachables. The default is **Enable**.

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

4.12 Loopback Interfaces

FASTPATH software provides for the creation, deletion, and management of loopback interfaces. They are

dynamic interfaces that are created and deleted via user-configuration. FASTPATH software supports multiple loopback interfaces.

A loopback interface is always expected to be up. As such, it provides a means to configure a stable IP address on the device that may be referred to by other switches. This interface provides the source address for sent packets and can receive both local and remote packets. It is typically used by routing protocols.

A loopback interface is a pseudo-device for assigning local addresses so that the router can be communicated with by this address, which is always up and can receive traffic from any of the existing active interfaces. Thus, given reachability from a remote client, the address of the loopback can be used to communicate with the router through various services such as telnet and SSH. In this way, the address on a loopback behaves identically to any of the local addresses of the router in terms of the processing of incoming packets.

The Routing > Loopbacks folder contains links to the following web pages that configure and display loopback parameters and data:

- Loopbacks Configuration
- Loopbacks Summary

4.12.1 Loopbacks Configuration

Use the Loopbacks Configuration page to create, configure, or remove loopback interfaces. You can also set up or delete a secondary address for a loopback.

To display the page, click **Routing > Loopbacks > Configuration** in the navigation tree. If no loopback interfaces exist on the system, the page only has two fields, as Figure 4-54 shows.

Figure 4-54: Loopback Configuration—Create

Additional fields display depending on whether or not a loopback has already been created, as shown in Figure 4-55.

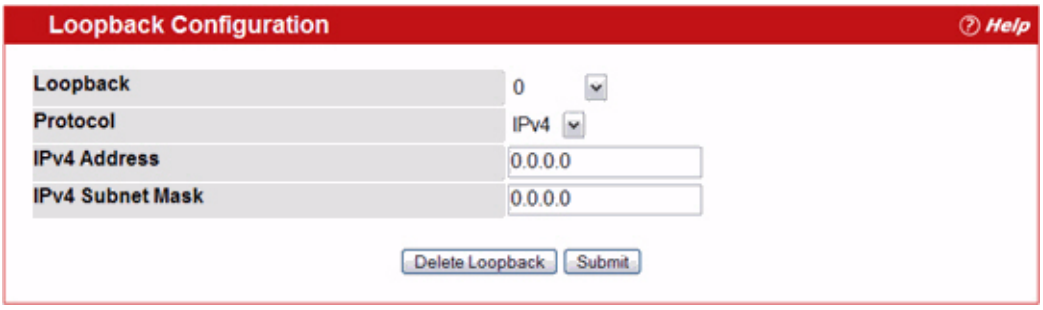

Figure 4-55: Configured Loopback Interface

The fields available on the Loopbacks Configuration page depend on whether any loopback interfaces exist and whether the protocol is IPv4 or IPv6. The following table describes all fields, which are not all on the

same screen at the same time.

Table 4-48: Configured Loopback Interface Fields

Field	Description
Loopback	Use the dropdown menu to select from the list of currently configured loopback interfaces. Create is also a valid choice if the maximum number of loopback interfaces has not been created.
Loopback ID	When Create is selected in the Loopback field, this list of available loopback IDs displays.
Protocol	Select IPv4 or IPv6 to configure the corresponding attributes on the loopback interface. The protocol selected affects the fields that are displayed on this page.
IPv6 Mode	Enable or disable IPv6 on this interface. This option only displays when the Protocol specified is IPv6, and is only configurable prior to specifying an explicit IPv6 address.
IPv6 Address	Select from the list of configured IPv6 addresses for the selected Loopback interface. Add is also a valid choice if the maximum number of addresses has not been configured. This option only displays when the Protocol specified is IPv6.
IPv6 Address	When Add is chosen from the IPv6 Address selector this IPv6 address input field becomes visible. Enter the address in the format of prefix/length. This option only displays when the Protocol specified is IPv6.
EU164	You also have the option to specify the 64-bit extended unique identifier (EUI-64). This option only displays when the Protocol specified is IPv6.
IPv4 Address	The primary IPv4 address for this interface in dotted decimal notation. This option only displays when the Protocol specified is IPv4.
IPv4 Subnet Mask	The primary IPv4 subnet mask for this interface in dotted decimal notation. This option only displays when the Protocol specified is IPv4.

The following fields display when a primary address is configured. You can configure multiple secondary addresses.

Table 4-49: Loopback Interface Secondary Address Fields

Field	Description
Secondary Address	Select a configured IPv4 secondary address for the selected Loopback interface from the dropdown menu. A new address can be entered in the Secondary IP Address field by selecting Add Secondary IP Address here (if the maximum number of secondary addresses has not been configured). A primary address must be configured before a secondary address can be added.
Secondary IP Address	The secondary IP address for this interface in dotted decimal notation. This input field is visible only when Add Secondary is selected.
Secondary Subnet Mask	The secondary subnet mask for this interface in dotted decimal notation. This input field is visible only when Add Secondary is selected.

4.12.1.1 Creating a New Loopback (IPv4)

1. From the **Loopbacks Configuration** page, select **Create** from the **Loopback** dropdown menu.
2. Specify an ID to use in the **Loopback ID** field.
3. Click **Submit**.

   The Loopback ID field goes away, and additional loopback fields display, as Figure 4-56 shows.

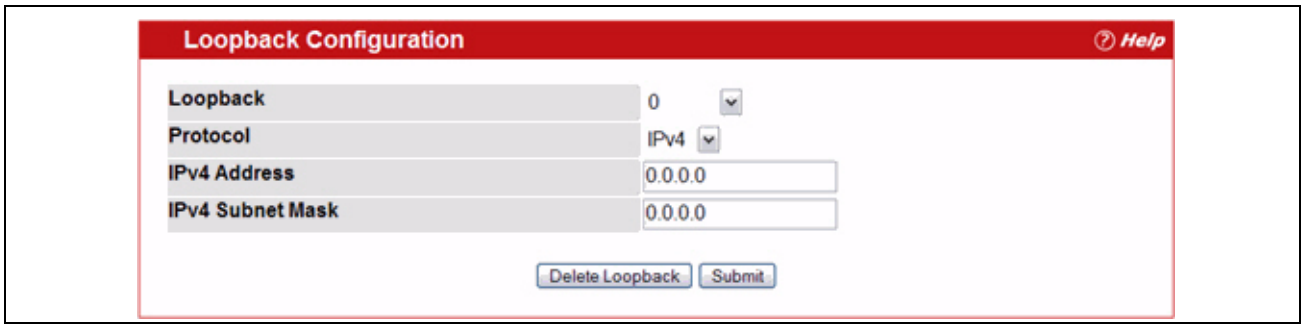

Figure 4-56: Loopbacks Configuration—IPv4 Entry

4. In the **Protocol** field, select **IPv4**
5. Enter desired values in the remaining fields.
6. Click **Submit**.

   The new loopback is saved, and the web page reappears showing secondary address configuration fields. For an example of the fields on this page, see Figure 4-55.
7. Optionally, complete the **Secondary Address**, **Secondary IP Address**, and **Secondary Subnet Mask** fields.
8. Click the **Add Secondary** button. The secondary address is saved, and the web page reappears showing the primary and secondary loopback addresses.

4.12.1.2 Creating a New Loopback (IPv6)

1. Open the **Loopbacks Configuration** page.
2. Select **Create** from the **Loopback** dropdown menu.
3. Specify an ID to use in the **Loopback ID** field.
4. Click **Submit**.

   The Loopback ID field goes away, and the remaining loopback fields display.
5. Choose **IPv6** from the drop down box in the **Protocol** field.

   The screen refreshes and displays the IPv6 loopback interface configuration fields, as Figure 4-57 shows.

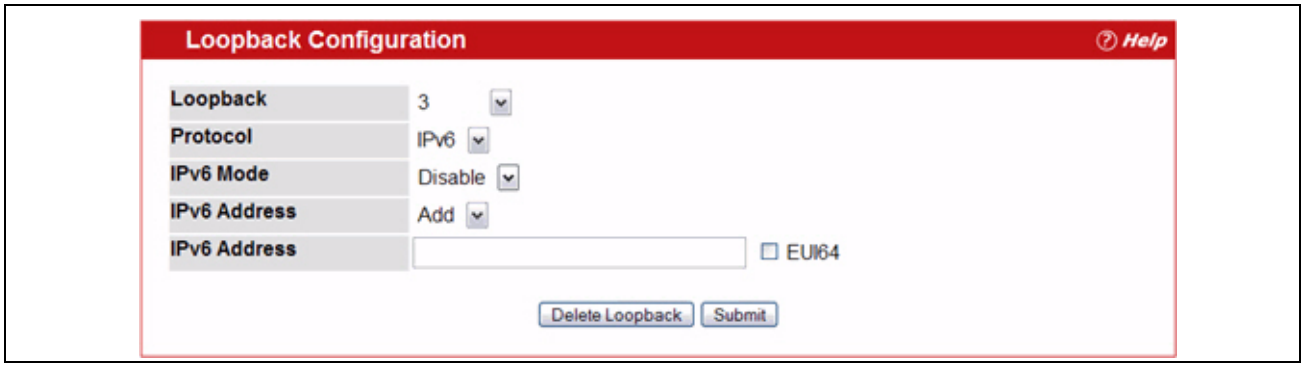

Figure 4-57: Loopbacks Configuration—IPv6 Entry

6. Add the **IPv6 Address**.
7. Enter desired values in the remaining fields.
8. Click **Submit**.

   The new loopback is saved, and the device is updated.

4.12.1.3 Configuring an Existing Loopback

1. Open the **Loopback Configuration** page.
2. Specify the loopback to configure in the **Loopback** dropdown menu.
3. Change field values as desired in the remaining fields.
4. Click **Apply Changes**.

   The new configuration is saved, and the device is updated.

4.12.1.4 Removing a Loopback

1. Open the **Loopback Configuration** page.
2. Specify the loopback to remove in the **Loopback** dropdown menu.
3. Click **Delete Loopback**.

   The loopback is deleted, and the device is updated.

4.12.1.5 Removing a Secondary Address

1. Open the **Loopback Configuration** page.
2. Specify the loopback to be affected.
3. Specify the secondary address to be removed.
4. Click **Delete Selected Secondary**.

   The secondary address is deleted, and the device is updated.

4.12.2 Loopbacks Summary

Use the Loopbacks Summary page to display a summary of configured loopbacks.

To display the page, click **Routing > Loopbacks > Summary** in the navigation tree.

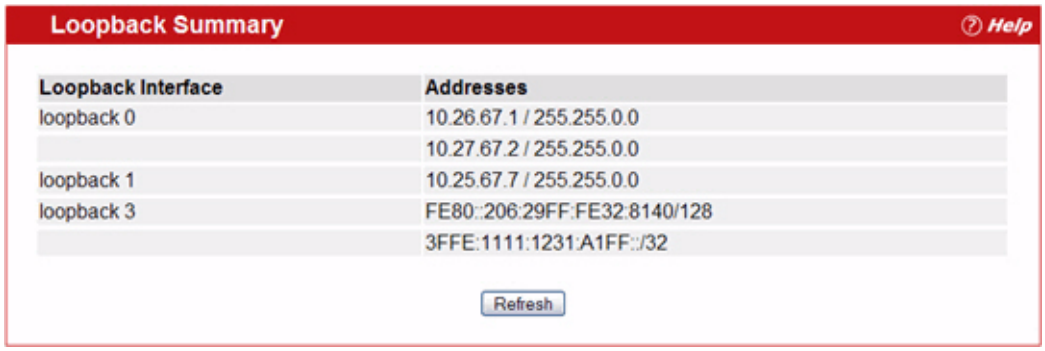

Figure 4-58: Loopbacks Summary

Table 4-50: Loopbacks Summary Fields

Field	Description
Loopback Interface	The ID of the configured loopback interface.
Addresses	A list of the addresses configured on the loopback interface.

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

5 Managing Device Security

Use the features in the Security folder on the navigation tree menu to set management security parameters for port, user, and server security.

The Security folder contains links to the following features:

- Port Access Control
- RADIUS Settings
- TACACS+ Settings
- Secure HTTP
- Secure Shell

5.1 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 the 802.1x features on the system:

- Global Port Access Control Configuration
- Port Configuration
- Port Status
- Port Summary
- Port Access Control Statistics
- Client Summary
- Client Detail
- Port Access Privileges
- Port Access Summary

5.1.1 Global Port Access Control Configuration

Use the Port Based Access Control Configuration page to enable or disable port access control on the system.

To display the Port Based Authentication page, click **Port Based Access Control > Configuration** in the navigation menu.

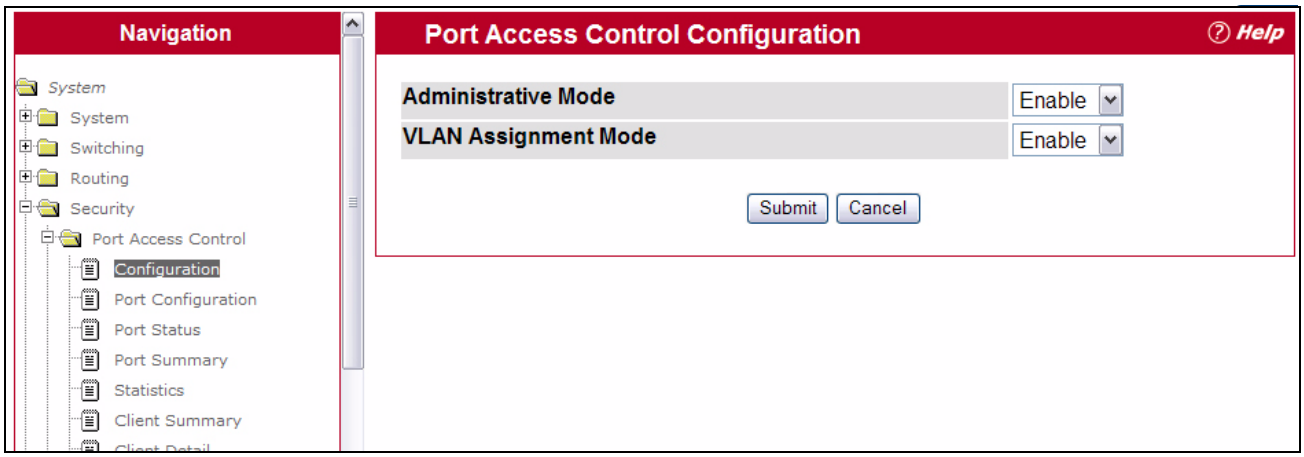

Figure 5-1: Port Access Control—Port Configuration

Table 5-1: Port Access Control—Port Configuration Fields

Field	Description
Administrative Mode	Select **Enable** or **Disable** 802.1x mode on the switch. The default is Disable. This feature permits port-based authentication on the switch.
VLAN Assignment Mode	If enabled, when a supplicant is authenticated by a authentication server, the port that the supplicant is connected to is placed in a particular VLAN specified by the RADIUS server. VLAN Assignment mode controls if the switch is allowed to place a port in a RADIUS-assigned VLAN. A port's VLAN assignment is determined by the first supplicant that is authenticated on the port.

If you change the mode, click **Submit** to apply the new settings to the system.

5.1.2 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 Based Access Control Port Configuration page, click **Security > Port Based Access Control > Port Configuration** in the navigation menu.

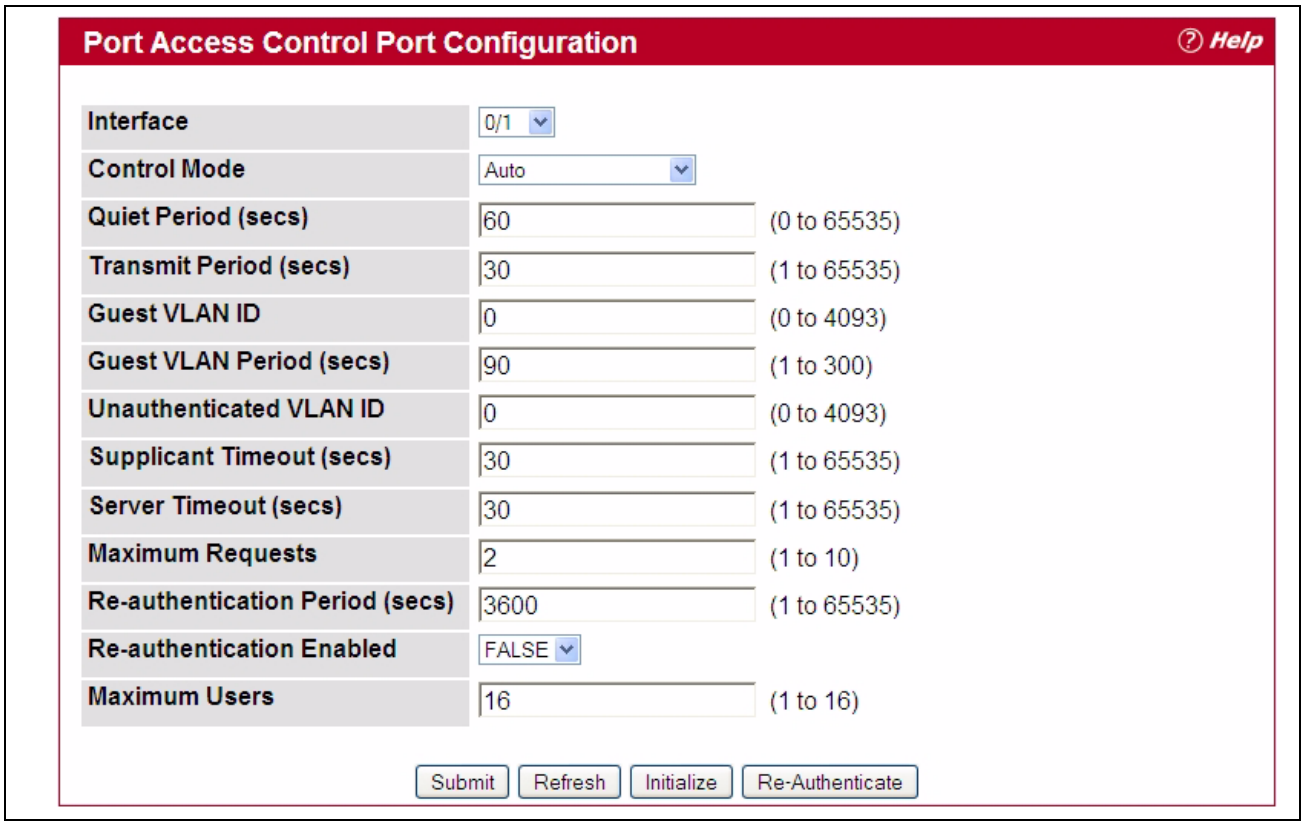

Figure 5-2: Port Access Control Port Configuration

Table 5-2: Port Access Control Port Configuration Fields

Field	Description
Port	Selects the Unit and Port to configure.
Control Mode	Defines the port authorization state. The control mode is only set if the link status of the port is link up. The possible field values are: • **Auto**: Automatically detects the mode of the interface. • **Force Authorized**: Places the interface into an authorized state without being authenticated. The interface sends and receives normal traffic without client port-based authentication. • **Force Unauthorized**: Denies the selected interface system access by moving the interface into unauthorized state. The switch cannot provide authentication services to the client through the interface. • **MAC-based**: Sets the mode of the interface to authentication on a per supplicant basis.
Quiet Period (secs)	Defines the amount of time that the switch remains in the quiet state following a failed authentication exchange. The possible field range is 0-65535. The field value is in seconds. The field default is 60 seconds.
Transmit Period (secs)	Defines the transmit period for the selected port. The transmit period is the value, in seconds, of the timer used by the authenticator state machine on the specified port to determine when to send an EAPOL EAP Request/Identity frame to the supplicant. The transmit period must be a number in the range of 1 and 65535. The default value is 30.
Guest VLAN ID	Defines the Guest VLAN ID on the interface. The valid range is 0 to 3965. The default value is 0. Changing the value will not change the configuration until you click the **Submit** button. Enter zero (0) to clear the Guest VLAN ID on the interface.

Table 5-2: Port Access Control Port Configuration Fields (Continued)

Field	Description
Guest VLAN Period (secs)	Defines the Guest VLAN period for the selected port. The Guest VLAN period is the value, in seconds, of the timer used by the Guest VLAN Authentication. The Guest VLAN timeout must be a value in the range of 1 to 300. The default value is 90. Changing the value will not change the configuration until you click the **Submit** button.
Unauthenticated VLAN ID	Defines the Unauthenticated VLAN ID for the selected port. The valid range is 0 to 3965. The default value is zero (0). Changing the value will not change the configuration until you click the **Submit** button. Enter zero (0) to clear the Unauthenticated VLAN ID on the interface.
Supplicant Timeout (secs)	Defines the amount of time that lapses before EAP requests are resent to the user. The value must be in the range of 1 to 65535 seconds. The value is 30 seconds. Changing the value will not change the configuration until you click the **Submit** button.
Server Timeout (secs)	Defines the amount of time that lapses before the switch resends a request to the authentication server. The field value is in seconds. The range is **1-65535**, and the field default is 30 seconds. Changing the value will not change the configuration until you click the **Submit** button.
Maximum Requests	Defines the maximum number of times the switch can send an EAP request before restarting the authentication process if it does not receive a response. The possible field range is 1-10. The field default is 2 retries.
Reauthentication Period (secs)	Indicates the time span in which the selected port is reauthenticated. The field value is in seconds. The range is 1 - 65535, and the field default is 3600 seconds. Changing the value will not change the configuration until you click the **Submit** button.
Reauthentication Enabled	Reauthenticates the selected port periodically, when enabled. The default value is False. Changing the value will not change the configuration until you click the **Submit** button.
Maximum Users	Defines the maximum number of clients that can get authenticated on the port in the MAC-based dot1x authentication mode. The range is 1 to 16. The default value is 16. Changing the value will not change the configuration until you click the **Submit** button.

- Click **Submit** to send the updated screen to the switch and cause the changes to take effect on the switch but these changes will not be retained across a power cycle unless a save is performed.
- Click **Refresh** to update the information on the screen.
- Click **Initialize** to begin the initialization sequence on the selected port. This button is only selectable if the control mode is 'auto'. If the button is not selectable, it will be grayed out. Once this button is pressed, the action is immediate. It is not required to press the Submit button for the action to occur.
- Click **Reauthenticate** to begin the reauthentication sequence on the selected port. This button is only selectable if the control mode is 'auto'. If the button is not selectable, it will be grayed out. Once this button is pressed, the action is immediate. It is not required to press the Submit button for the action to occur.

5.1.3 Port Status

Use the Port Access Control Port Status page to view information about the port access control settings on a specific port.

To access the Port Based Access Control Port Status page, click **Security > Port Based Access Control > Port Status** in the navigation menu.

Port Access Control Port Status Help

Interface	0/1
Protocol Version	Version1
PAE Capabilities	Authenticator
Control Mode	Auto
Authenticator PAE State	Initialize
Backend State	Initialize
Quiet Period (secs)	60
Transmit Period (secs)	30
Guest VLAN ID	0
Guest VLAN Period (secs)	90
Supplicant Timeout (secs)	30
Server Timeout (secs)	30
Maximum Requests	2
VLAN Assigned	0
VLAN Assigned Reason	Not Assigned
Reauthentication Period (secs)	3600
Reauthentication Enabled	FALSE
Key Transmission Enabled	FALSE
Control Direction	Both
Maximum Users	16
Unauthenticated VLAN ID	0
Session Timeout	0
Session Termination Action	Default

Refresh

Figure 5-3: Port Access Control Status

Figure 5-4 on page 335 is an example of the fields displayed for the port when the Control mode of the port is MAC-based.

Port Access Control Port Status Help

Interface	0/1
Protocol Version	Version1
PAE Capabilities	Authenticator
Control Mode	MAC Based
Quiet Period (secs)	60
Transmit Period (secs)	30
Guest VLAN ID	0
Guest VLAN Period (secs)	90
Supplicant Timeout (secs)	30
Server Timeout (secs)	30
Maximum Requests	2
Reauthentication Period (secs)	3600
Reauthentication Enabled	FALSE
Key Transmission Enabled	FALSE
Control Direction	Both
Maximum Users	16
Unauthenticated VLAN ID	0

Logical Port	Supplicant MAC Address	Authenticator PAE State	Backend Authentication State	VLAN Assigned	VLAN Assigned Reason

Refresh

Figure 5-4: Port Access Control Status - MAC-based Control Mode

Table 5-3: Port Access Control Status Fields

Field	Description
Port	Selects the Unit and Port to view.
Protocol Version	This field displays the protocol version associated with the selected port. The only possible value is 1, corresponding to the first version of the 802.1x specification. This field is not configurable.
PAE Capabilities	This field displays the port access entity (PAE) functionality of the selected port. Possible values are "Authenticator" or "Supplicant". This field is not configurable.
Control Mode	Defines the port authorization state. The control mode is only set if the link status of the port is link up. The possible field values are: • **Auto**: Automatically detects the mode of the interface. • **Force Authorized**: Places the interface into an authorized state without being authenticated. The interface sends and receives normal traffic without client port-based authentication. • **Force Unauthorized**: Denies the selected interface system access by moving the interface into unauthorized state. The switch cannot provide authentication services to the client through the interface. • **MAC-based**: Sets the mode of the interface to authentication on a per supplicant basis.
Authenticator PAE State	This field displays the current state of the authenticator PAE state machine. Possible values are as follows: • Initialize • Disconnected • Connecting • Authenticating • Authenticated • Aborting • Held • ForceAuthorized • ForceUnauthorized
Backend Authentication State	This field displays the current state of the backend authentication state machine. Possible values are as follows: • Request • Response • Success • Fail • Timeout • Initialize • Idle
Quiet Period	Displays the configured quiet period for the selected port. This quiet period is the value, in seconds, of the timer used by the authenticator state machine on this port to define periods of time in which it will not attempt to acquire a supplicant. The quiet period is the period for which the authenticator does not attempt to acquire a supplicant after a failed authentication exchange with the supplicant. The quiet period is a number in the range of 0 and 65535.
Transmit Period	Displays the configured transmit period for the selected port. The transmit period is the value, in seconds, of the timer used by the authenticator state machine on the specified port to determine when to send an EAPOL EAP Request/Identity frame to the supplicant. The transmit period is a number in the range of 1 and 65535.
Guest VLAN ID	Displays the Guest VLAN ID configured on the interface. The valid range is 0 to 3965.
Guest VLAN Period (secs)	Displays the Guest VLAN period for the selected port. The Guest VLAN period is the value, in seconds, of the timer used by the Guest VLAN Authentication. The value is in the range of 1 to 300.

Table 5-3: Port Access Control Status Fields (Continued)

Field	Description
Supplicant Timeout	Displays the configured supplicant timeout for the selected port. The supplicant timeout is the value, in seconds, of the timer used by the authenticator state machine on this port to timeout the supplicant. The supplicant timeout is a value in the range of 1 and 65535.
Server Timeout	Displays the configured server timeout for the selected port. The server timeout is the value, in seconds, of the timer used by the authenticator on this port to timeout the authentication server. The server timeout is a value in the range of 1 and 65535.
Maximum Requests	Displays the configured maximum requests for the selected port. The maximum requests value is the maximum number of times the authenticator state machine on this port will retransmit an EAPOL EAP Request/Identity before timing out the supplicant. The maximum requests value is in the range of 1 and 10.
VLAN Assigned	Displays the VLAN ID assigned to the selected interface by the Authenticator. **Note:** This field is displayed only when the port control mode of the selected interface is not MAC-based.
VLAN Assigned Reason	Displays the reason for the VLAN ID assigned by the authenticator to the selected interface. Possible values are: • Radius • Unauth • Default • Not Assigned **Note:** This field is displayed only when the port control mode of the selected interface is not MAC-based.
Reauthentication Period	Displays the configured reauthentication period for the selected port. The reauthentication period is the value, in seconds, of the timer used by the authenticator state machine on this port to determine when reauthentication of the supplicant takes place. The reauthentication period is a value in the range of 1 and 65535.
Reauthentication Enabled	Displays if reauthentication is enabled on the selected port. This is a configurable field. The possible values are 'true' and 'false'. If the value is 'true' reauthentication will occur. Otherwise, reauthentication will not be allowed.
Key Transmission Enabled	This field displays if key transmission is enabled on the selected port. This is not a configurable field. The possible values are 'true' and 'false'. If the value is 'false', key transmission will not occur. Otherwise, key transmission is supported on the selected port.
Control Direction	This displays the control direction for the specified port. The control direction dictates the degree to which protocol exchanges take place between Supplicant and Authenticator. This affects whether the unauthorized controlled port exerts control over communication in both directions (disabling both incoming and outgoing frames) or just in the incoming direction (disabling only the reception of incoming frames). **Note:** This field is not configurable on some platforms.
Maximum Users	Displays the maximum number of clients that can get authenticated on the port in the MAC-based dot1x authentication mode. This field is configurable. The maximum users value is in range of 1 to 16.
Unauthenticated VLAN ID	Displays the Unauthenticated VLAN ID for the selected port. The valid range is 0 to 3965.

Table 5-3: Port Access Control Status Fields (Continued)

Field	Description
Session Timeout	Displays the Session Timeout set by the RADIUS Server for the selected port. **Note:** This field is displayed only when the port control mode of the selected port is not MAC-based.
Session Termination Action	Displays the Termination Action set by the RADIUS Server for the selected port. Possible values are: • Default • Reauthenticate If the termination action is Default then, at the end of the session, the client details are initialized. Otherwise, re-authentication is attempted. **Note:** This field is displayed only when the port control mode of the selected port is not MAC-based.
Logical Port	Displays the logical port number associated with the supplicant that is connected to the port. This field is not configurable. **Note:** This field is displayed when the port control mode of the selected port is MAC-based.
Supplicant MacAddress	This field displays the supplicant's MAC address that is connected to the port. This field is not configurable. **Note:** This field is displayed when the port control mode of the selected port is MAC-based.

5.1.4 Port Summary

Use the Port Access Control Port Summary page to view summary information about the port access control settings on all physical ports.

To access the Port Based Access Control Port Summary page, click **Security > Port Based Access Control > Port Summary** in the navigation menu.

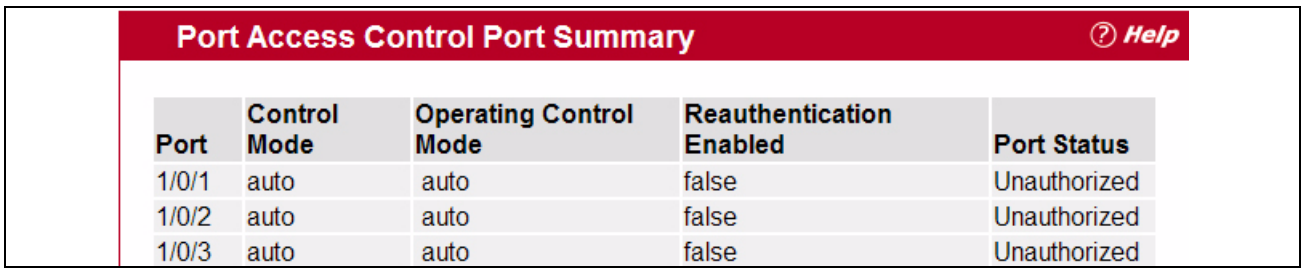

Port	Control Mode	Operating Control Mode	Reauthentication Enabled	Port Status
1/0/1	auto	auto	false	Unauthorized
1/0/2	auto	auto	false	Unauthorized
1/0/3	auto	auto	false	Unauthorized

Figure 5-5: Port Access Control Port Summary

Table 5-4: Port Access Control Port Summary Fields

Field	Description
Port	Selects the Unit and Port to view.
Control Mode	Displays the port authorization state. The possible field values are: • **Auto**: Automatically detects the mode of the interface. • **Force Authorized**: Places the interface into an authorized state without being authenticated. The interface sends and receives normal traffic without client port-based authentication. • **Force Unauthorized**: Denies the selected interface system access by moving the interface into unauthorized state. The switch cannot provide authentication services to the client through the interface. • **MAC-based**: Sets the mode of the interface to authentication on a per supplicant basis.
Operating Control Mode	Indicates the control mode under which the port is actually operating. Possible values are as follows: • ForceUnauthorized • ForceAuthorized • Auto • MAC-based • N/A: If the port is in detached state it cannot participate in port access control.
Reauthentication Enabled	Displays whether reauthentication is enabled on the port. This is a configurable field. The possible values are as follows: • **True**: Reauthentication will occur. • **False**: Reauthentication will not be allowed.
Port Status	Shows the authorization status of the port, which might be Authorized, Unauthorized or N/A. The value is N/A if the port is in detached state and cannot participate in port access control.

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

5.1.5 Port Access Control Statistics

Use the Port Access Control Statistics page to view EAP and EAPOL information on a specific port.

To access the Port Based Access Control Statistics page, click **Security > Port Based Access Control > Statistics** in the navigation menu.

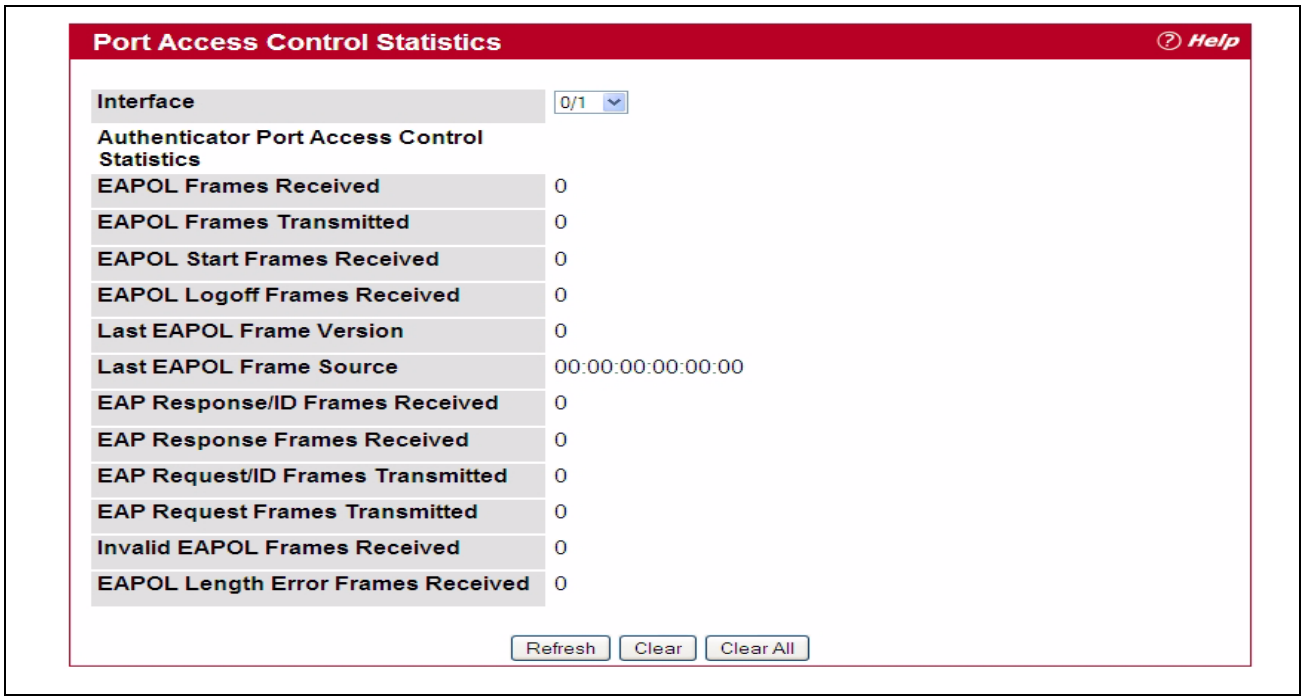

Figure 5-6: Port Access Control Statistics

Table 5-5: Port Access Control Statistics Fields

Field	Description
Port	Selects the port to be displayed. When the selection is changed, a screen refresh will occur causing all fields to be updated for the newly selected port. All physical interfaces are valid.
EAPOL Frames Received	Displays the number of valid EAPOL frames received on the port.
EAPOL Frames Transmitted	Displays the number of EAPOL frames transmitted via the port.
EAPOL Start Frames Received	Displays the number of EAPOL Start frames received on the port.
EAPOL Logoff Frames Received	Displays the number of EAPOL Log off frames that have been received on the port.
Last EAPOL Frames Version	Displays the protocol version number attached to the most recently received EAPOL frame.
Last EAPOL Frames Source	Displays the source MAC Address attached to the most recently received EAPOL frame.
EAP Response/ID Frames Received	Displays the number of EAP Respond ID frames that have been received on the port.
EAP Response Frames Received	Displays the number of valid EAP Respond frames received on the port.
EAP Request/ID Frames Transmitted	Displays the number of EAP Requested ID frames transmitted via the port.
EAP Request Frames Transmitted	Displays the number of EAP Request frames transmitted via the port.
Invalid EAPOL Frames Received	Displays the number of unrecognized EAPOL frames received on this port.
EAPOL Length Error Frames Received	Displays the number of EAPOL frames with an invalid Packet Body Length received on this port.

- Click **Refresh** to update the information on the page.
- Click **Clear All** to reset all statistics for all ports to 0. There is no confirmation prompt. When you click this button, the statistics are immediately cleared.
- Click **Clear** to reset the statistics for the selected port. There is no confirmation prompt. When you click this button, the statistics are immediately cleared.

5.1.6 Client Summary

Use the Port Access Control Client Summary page to view summary information about the supplicant device.

To access the Port Access Control Client Summary page, click **Security > Port Access Control > Client Summary** in the navigation menu.

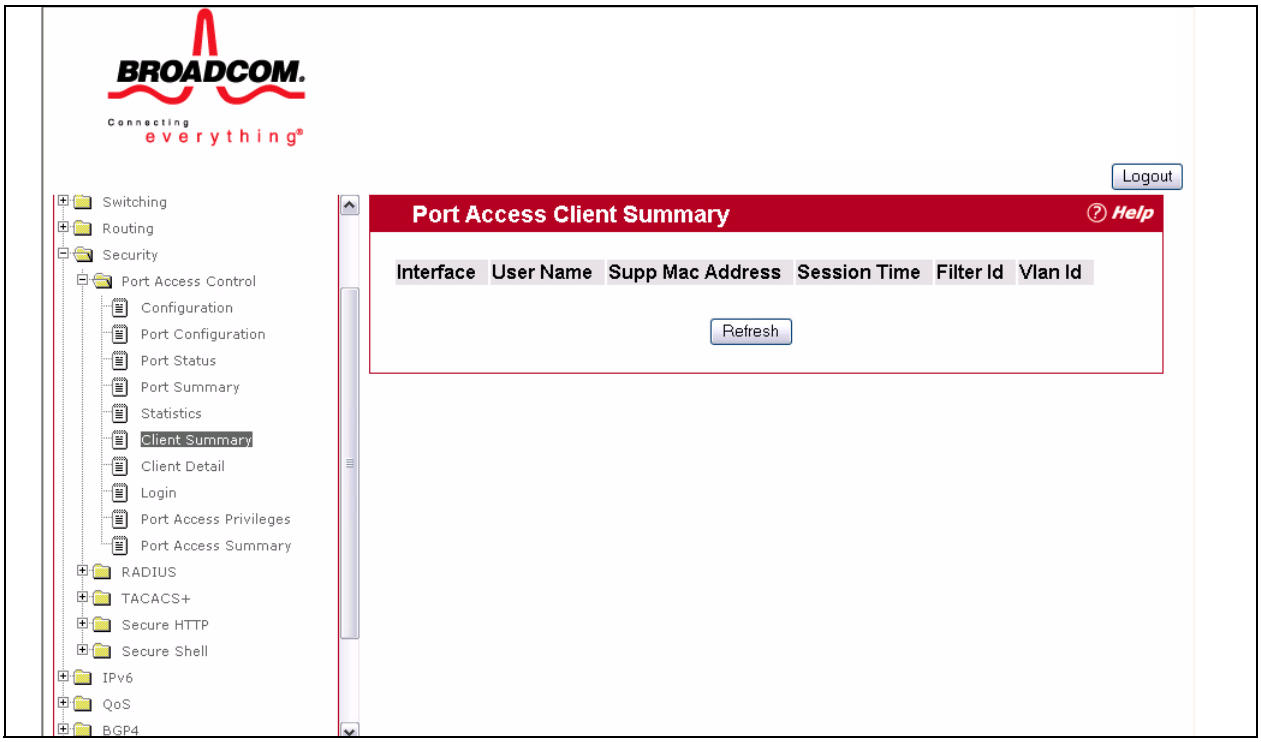

Figure 5-7: Port Access Control Client Summary

Table 5-6: Port Access Control Client Summary Fields

Field	Description
Interface	Displays the interface address of the supplicant device.
User Name	Displays the user name representing the supplicant device.
Supp Mac Address	Displays the supplicant device's MAC address.
Session Time	Displays the time since the supplicant logged in. The value is in seconds.
Filter ID	The policy filter ID assigned by the authenticator to the supplicant device.
VLAN ID	The VLAN ID assigned by the authenticator to the supplicant device.

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

5.1.7 Client Detail

Use the Port Access Control Client Detail page to view detail information about the supplicant device.

To access the Port Access Control Client Detail page, click **Security > Port Access Control > Client Detail** in the navigation menu.

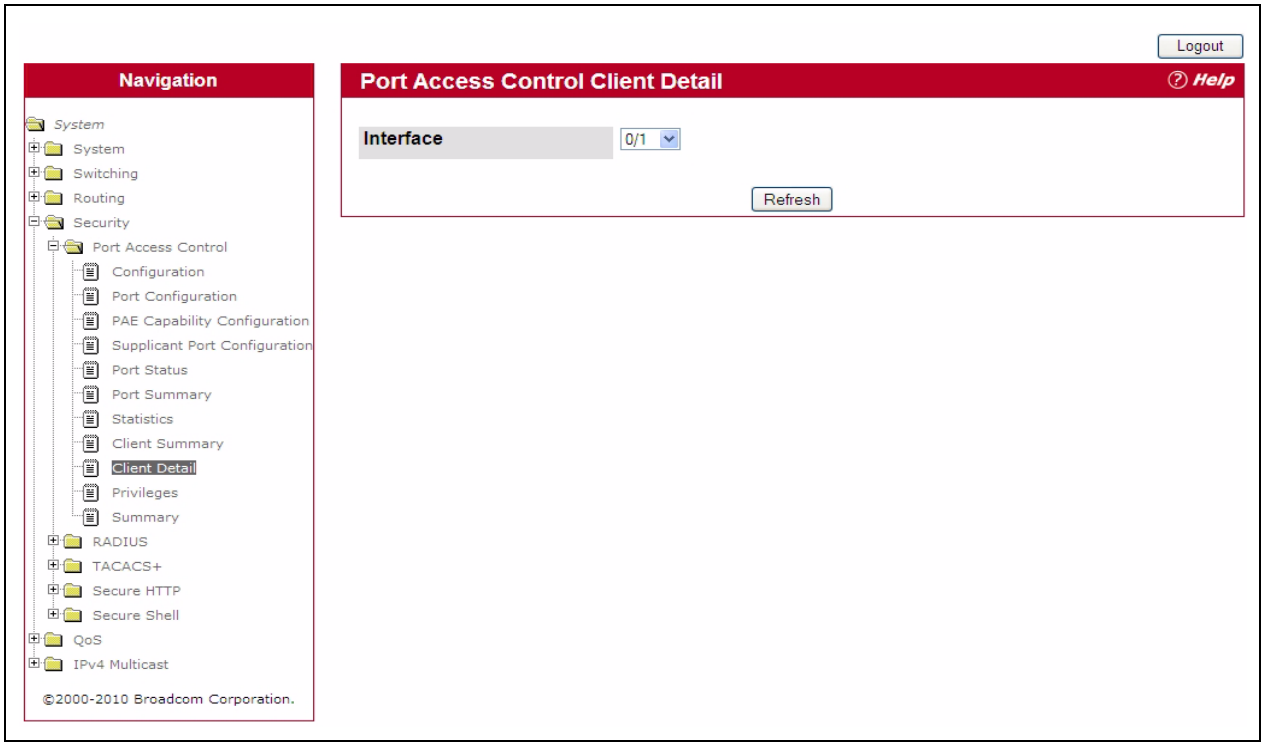

Figure 5-8: Port Access Control Client Detail

Table 5-7: Port Access Control Client Detail Fields

Field	Description
Port	Select the port address of the supplicant device.
User Name	Displays the user name representing the supplicant device.
Supplicant MAC Address	Displays the supplicant device's MAC address.
Session Time	Displays the time since the supplicant logged in. The value is in seconds.
Filter ID	The policy filter ID assigned by the authenticator to the supplicant device.
VLAN ID	The VLAN ID assigned by the authenticator to the supplicant device.
VLAN Assigned	Displays the reason for the VLAN ID assigned by the authenticator to the supplicant device.
Session Timeout	Displays the session timeout set by the radius server to the supplicant device.
Termination Action	Displays the termination action set by the radius server to the supplicant device.

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

5.1.8 Port Access Privileges

Use the Port Access Control Privileges page to grant or deny port access to users configured on the system.

To access the Port Based Access Control Privileges page, click **Security > Port Based Access Control > Privileges** in the navigation menu.

Figure 5-9: Port Access Control Privileges

Table 5-8: Port Access Privileges Fields

Field	Description
Port	Selects the port to grant or deny access. To grant or deny port access privileges to a user on all ports, select All from the drop-down menu.
Users	Lists the users configured on the system. The users that are highlighted have access to the selected port. By default, all users have access to all ports. To deny access to a port, Shift + click to select only the users to allow access. Make sure the username to deny port access is not selected, and then click **Submit**.

5.1.9 Port Access Summary

Use the Port Access Control Summary page to view a summary of which users are allowed access to the physical ports on the system.

To access the Port Based Access Control Summary page, click **Security > Port Based Access Control > Summary** in the navigation menu.

Port Access Summary Help

Interface	Users
0/1	admin guest
0/2	admin guest
0/3	admin guest
0/4	admin guest
0/5	admin guest
0/6	admin guest

Figure 5-10: Port Access Control Summary

Table 5-9: Port Access Summary Fields

Field	Description
Port	Lists the physical ports on the system.
Users	Lists the users that are allowed 802.1x access to the port. If a username is configured on the system and does not appear in the **Users** column for a port, the user is denied access to the port.

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

5.2 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
- Access Control Port (802.1x)

The RADIUS folder contains links to the following pages that help you view and configure system RADIUS settings:

- RADIUS Configuration
- Server Configuration
- Server Statistics
- Accounting Server Configuration

- Accounting Server Statistics
- Clear Statistics

RADIUS Configuration

Use the RADIUS Configuration page to view and configure various settings for the RADIUS servers configured on the system.

To access the RADIUS **Configuration** page, click **Security > RADIUS > Configuration** in the navigation menu.

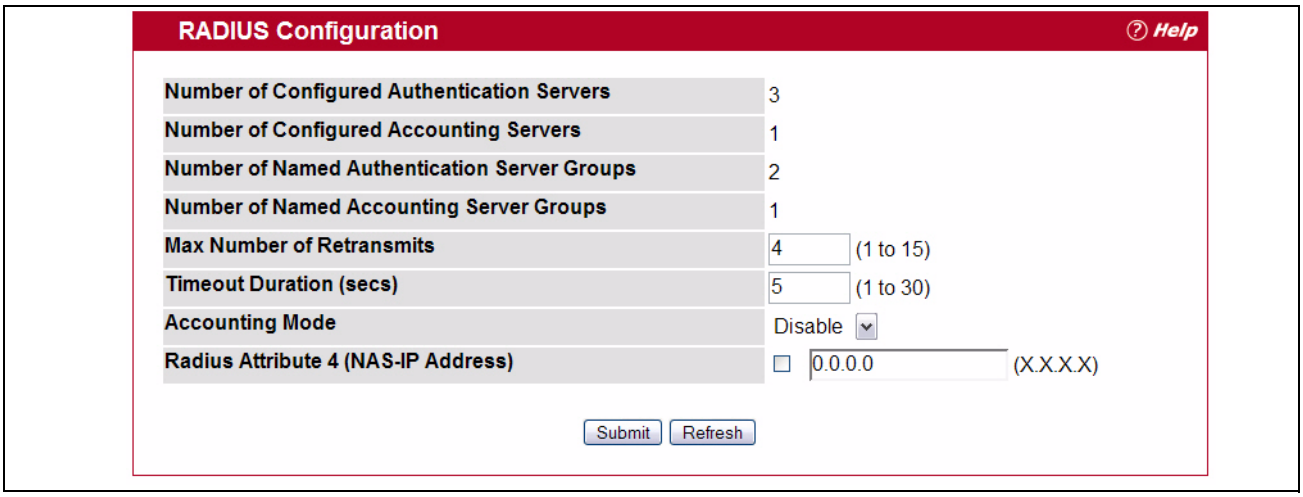

Figure 5-11: RADIUS Configuration

Table 5-10: RADIUS Configuration Fields

Field	Description
Number of Configured Authentication Servers	The number of RADIUS authentication servers configured on the system. The value can range from 0 to 32.
Number of Configured Accounting Servers	The number of RADIUS accounting servers configured on the system. The value can range from 0 to 32.
Number of Named Authentication Server Groups	The number of authentication server groups configured on the system. An authentication server group contains one or more configured authentication servers that share the same RADIUS server name.
Number of Named Accounting Server Groups	The number of accounting server groups configured on the system. An accounting server group contains one or more configured authentication servers that share the same RADIUS server name.

Table 5-10: RADIUS Configuration Fields (Continued)

Field	Description
Max Number of Retransmits	The value of the maximum number of times a request packet is retransmitted. The valid range is 1-15. Consideration to maximum delay time should be given when configuring RADIUS max retransmit and RADIUS timeout. If multiple RADIUS servers are configured, the max retransmit value on each will be exhausted 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 application equals the sum of (retransmit times 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 (secs)	The timeout value, in seconds, for request retransmissions. The valid range is 1 - 30. See the Max Number of Retransmits field description for more information about configuring the timeout duration.
Accounting Mode	Use the menu to select whether the RADIUS accounting mode is enabled or disabled on the current server.
RADIUS Attribute 4 (NAS-IP Address)	To set the network access server (NAS) IP address for the RADIUS server, select the option and enter the IP address of the NAS in the available field. The address should be unique to the NAS within the scope of the RADIUS server. The NAS IP address is only used in Access-Request packets.

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

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

5.2.1 Server Configuration

From the **Server Configuration** page, you can add a new RADIUS server, configure settings for a new or existing RADIUS server, and view RADIUS server status information. The RADIUS client on the switch supports up to 32 named authentication and accounting servers.

To access the RADIUS Server Configuration page, click **Security > R**AD**IUS > Server Configuration** in the navigation menu.

If there are no RADIUS servers configured on the system or if you select Add from the RADIUS Server Host Address menu, the fields described in the following table are available.

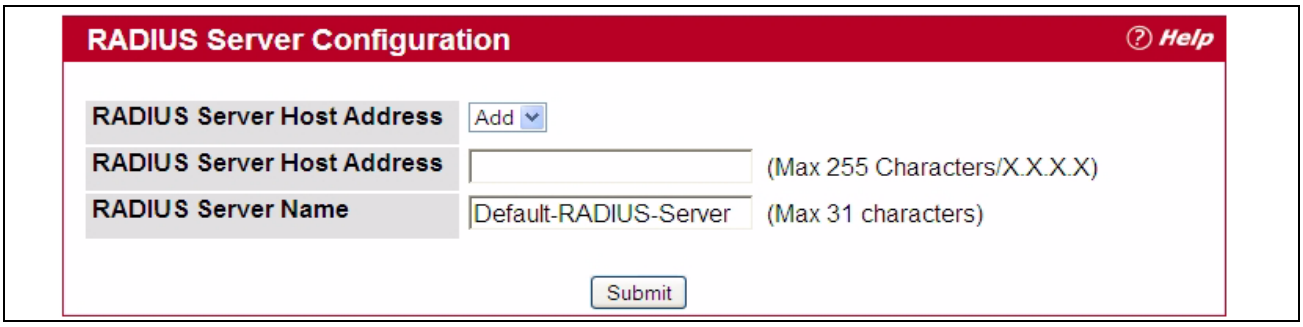

Figure 5-12: RADIUS Server Configuration—Add Server

Table 5-11: RADIUS Server Configuration Fields

Field	Description
RADIUS Server Host Address	To configure a new RADIUS server, select the Add option from the menu. To view or configure a RADIUS server that is already configured on the system, select its IP address from the menu.
Host Address	Enter the IP address of the RADIUS server to add. This field is only available when Add is selected in the **RADIUS Server Host Address** field.
RADIUS Server Name	Enter the name of the RADIUS server. The name can contain up to 32 alphanumeric characters. Spaces, hyphens, and underscores are also permitted. If you do not assign a name, the server is assigned the default name Default-RADIUS-Server. You can use the same name for multiple RADIUS Authentication servers. RADIUS clients can use RADIUS servers with the same name as backups for each other.

After you enter the RADIUS server information, click **Submit** to apply the changes to the system. The page refreshes, and additional RADIUS server configuration fields appear.

If at least one RADIUS server is configured on the switch, and a host address is selected in the RADIUS Server Host Address field, then additional fields are available on the RADIUS Server Configuration page. After you add a RADIUS server, use the Server Configuration page to configure the server settings.

If you select **Add** from the RADIUS Server Host Address field, the page refreshes and several of the configuration options are hidden.

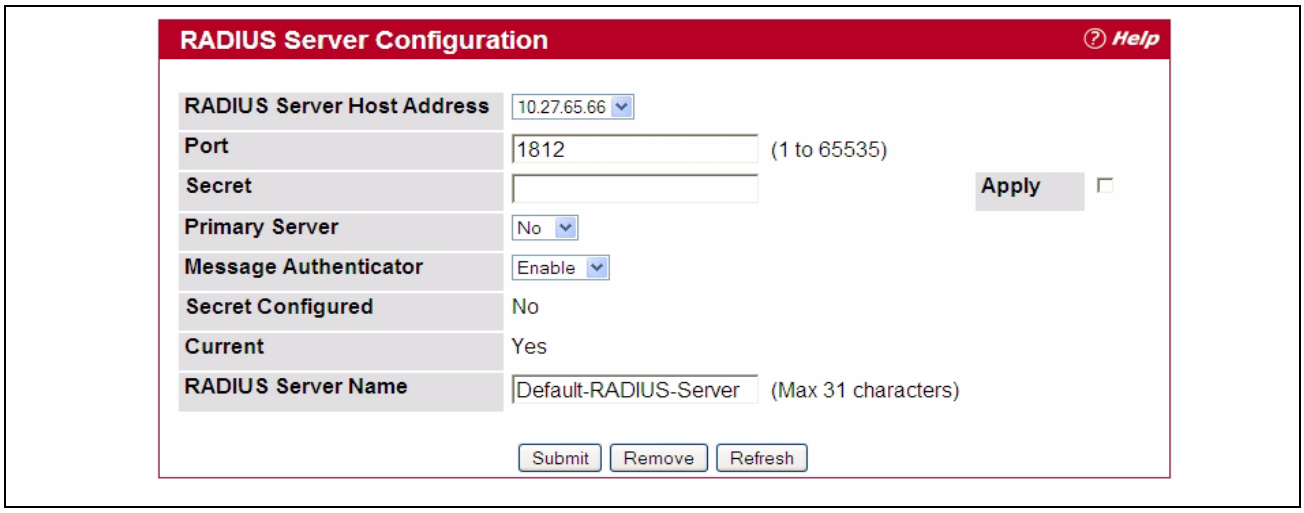

Figure 5-13: RADIUS Server Configuration—Server Added

Table 5-12: RADIUS Server Configuration Fields

Field	Description
RADIUS Server Host Address	Use the drop-down menu to select the IP address of the RADIUS server to view or configure. Select Add to configure additional RADIUS servers.
Port	Identifies the authentication port the server uses to verify the RADIUS server authentication. The port is a UDP port, and the valid range is 1-65535. The default port for RADIUS authentication is 1812.
Secret	Shared secret text string used for authenticating and encrypting all RADIUS communications between the device and the RADIUS server. This secret must match the RADIUS encryption.
Apply	The Secret will only be applied if this box is checked. If the box is not checked, anything entered in the Secret field will have no affect and will not be retained. This field is only displayed if the user has READWRITE access.
Primary Server	Sets the selected server to the Primary (Yes) or Secondary (No) server. If you configure multiple RADIUS servers with the same RAIDUS Server Name, designate one server as the primary and the other(s) as the backup server(s). The switch attempts to use the primary server first, and if the primary server does not respond, the switch attempts to use one of the backup servers with the same RADIUS Server Name.

Table 5-12: RADIUS Server Configuration Fields (Continued)

Field	Description
Message Authenticator	Enable or disable the message authenticator attribute for the selected server.
Secret Configured	Indicates whether the shared secret for this server has been configured.
Current	Indicates whether the selected RADIUS server is the current server (Yes) or a backup server (No). 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. If the primary server is not configured, the current server is the most recently configured RADIUS server.
RADIUS Server Name	Shows the RADIUS server name. To change the name, enter up to 32 alphanumeric characters. Spaces, hyphens, and underscores are also permitted. If you do not assign a name, the server is assigned the default name Default-RADIUS-Server. You can use the same name for multiple RADIUS Authentication servers. RADIUS clients can use RADIUS servers with the same name as backups for each other.

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.

To delete a configured RADIUS authentication server, select the IP address of the server from the **RADIUS Server Host Address** menu, and then click **Remove**.

- Click **Refresh** to update the page with the most current information.

5.2.1.1 Named Server Status Information

The RADIUS Named Server Status page shows summary information about the RADIUS servers configured on the system.

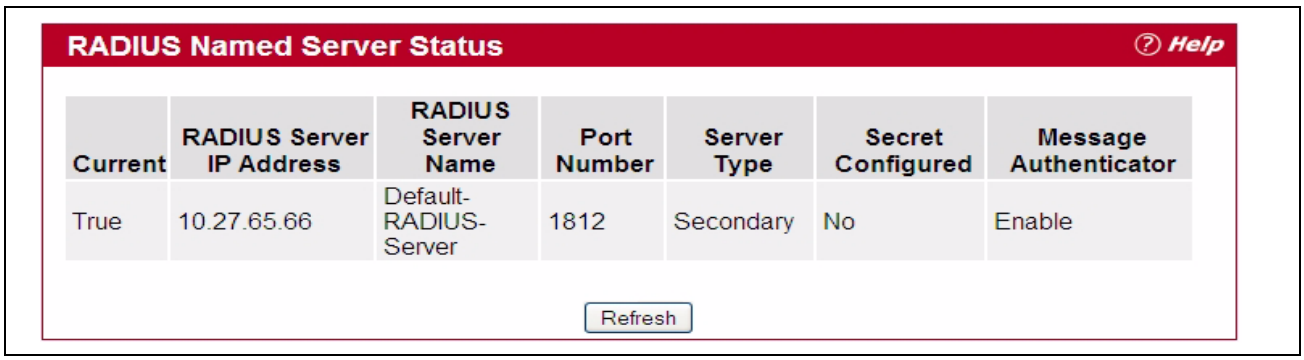

Figure 5-14: Named Server Status

Table 5-13: RADIUS Server Configuration Fields

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

Click **Refresh** to update the page with the most current information.

5.2.2 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 > Server Statistics** in the navigation menu.

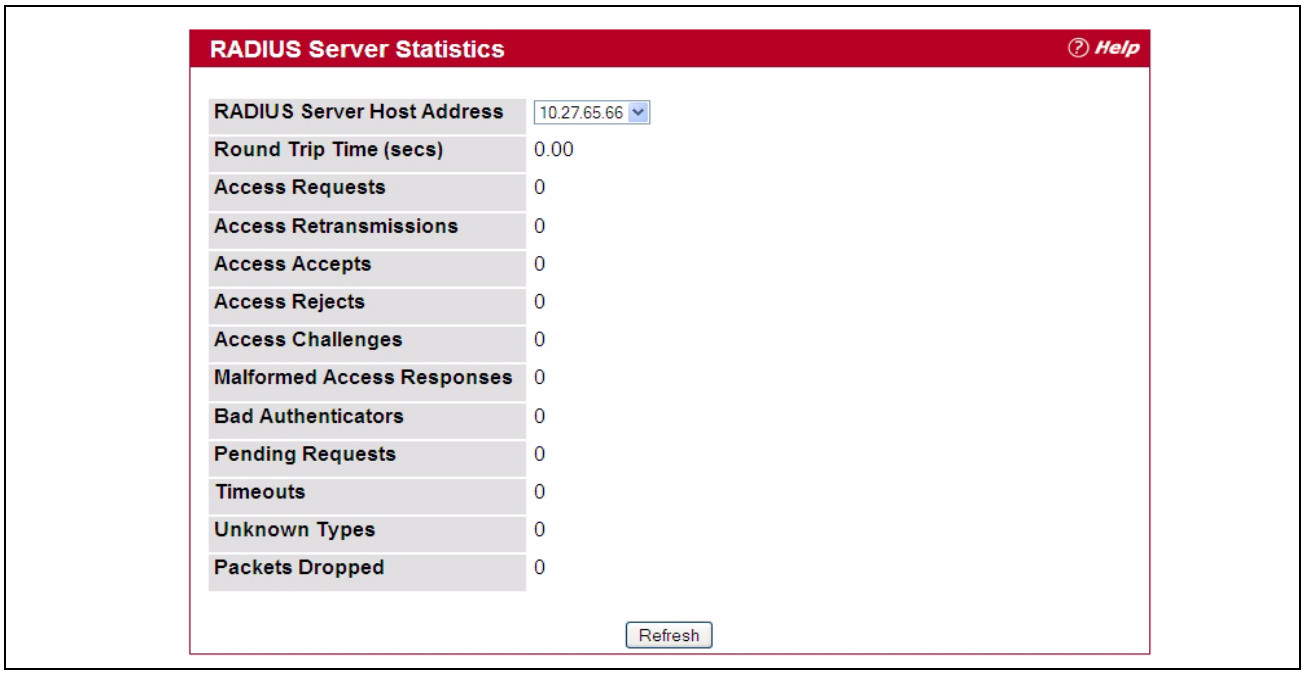

Figure 5-15: RADIUS Server Statistics

Table 5-14: RADIUS Server Statistics Fields

Field	Description
RADIUS Server Host Address	Use the drop-down menu to select the IP address of the RADIUS server for which to display statistics.
Round Trip Time (secs)	The time interval, in hundredths of a second, between the most recent Access-Reply/Access-Challenge and the Access-Request that matched it from this RADIUS authentication server.
Access Requests	The number of RADIUS Access-Request packets sent to this server. This number does not include retransmissions.
Access Retransmissions	The number of RADIUS Access-Request packets retransmitted to this server.
Access Accepts	The number of RADIUS Access-Accept packets, including both valid and invalid packets, that were received from this server.
Access Rejects	The number of RADIUS Access-Reject packets, including both valid and invalid packets, that were received from this server.
Access Challenges	The number of RADIUS Access-Challenge packets, including both valid and invalid packets, that were received from this server.
Malformed Access Responses	The number of malformed RADIUS Access-Response packets received from this server. Malformed packets include packets with an invalid length. Bad authenticators or signature attributes or 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 this server.

Table 5-14: RADIUS Server Statistics Fields (Continued)

Field	Description
Pending Requests	The number of RADIUS Access-Request packets destined for this server that have not yet timed out or received a response.
Timeouts	The number of authentication timeouts to this server.
Unknown Types	The number of RADIUS packets of unknown type which were received from this server on the authentication port.
Packets Dropped	The number of RADIUS packets received from this server on the authentication port and dropped for some other reason.

Click **Refresh** to update the page with the most current information.

5.2.3 Accounting Server Configuration

From the **Accounting Server Configuration** page, you can add a new RADIUS accounting server, configure settings for a new or existing RADIUS accounting server, and view RADIUS accounting server status information. The RADIUS client on the switch supports up to 32 named authentication and accounting servers.

If there are no RADIUS accounting servers configured on the system or if you select Add from the Accounting Server Host Address menu, the fields described in the following table are available.

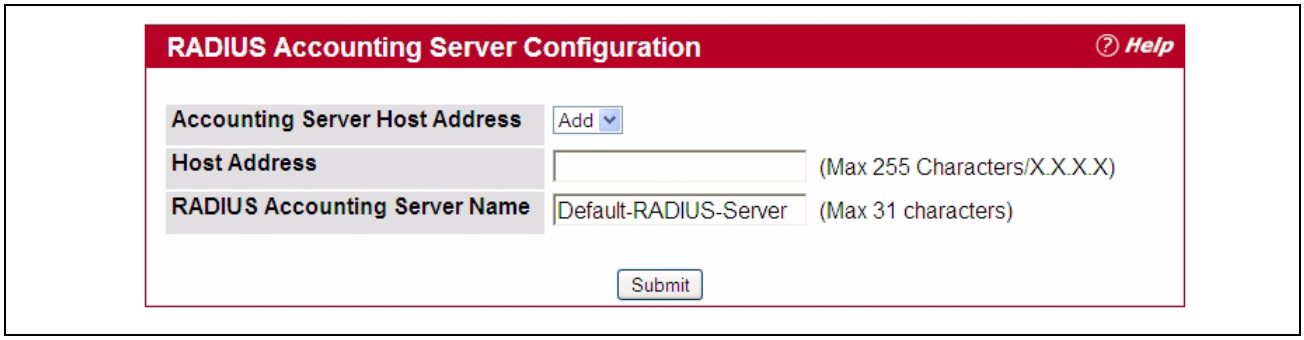

Figure 5-16: Add RADIUS Accounting Server

Table 5-15: RADIUS Server Configuration Fields

Field	Description
Accounting Server Host Address	To configure a new RADIUS accounting server, select the Add option from the menu. To view or configure an accounting server that is already configured on the system, select its IP address from the menu.
Host Address	Enter the IP address of the RADIUS accounting server to add. This field is only available when Add is selected in the **Accounting Server Host Address** field.
RADIUS Accounting Server Name	Enter a name for the RADIUS accounting server. The name can contain up to 32 alphanumeric characters. Spaces, hyphens, and underscores are also permitted. If you do not assign a name, the server is assigned the default name Default-RADIUS-Server. You can use the same name for multiple RADIUS accounting servers. RADIUS clients can use accounting servers with the same name as backups for each other.

After you enter the RADIUS accounting server information, click **Submit** to apply the changes to the system.The page refreshes, and additional accounting server configuration fields appear.

If at least one RADIUS accounting server is configured on the switch, and a host address is selected in the Accounting Server Host Address field, then additional fields are available on the Accounting Server Configuration page. After you add an accounting server, use the Accounting Server Configuration page to configure the server settings.

If you select **Add** from the Accounting Server Host Address field, the page refreshes and several of the configuration options are hidden.

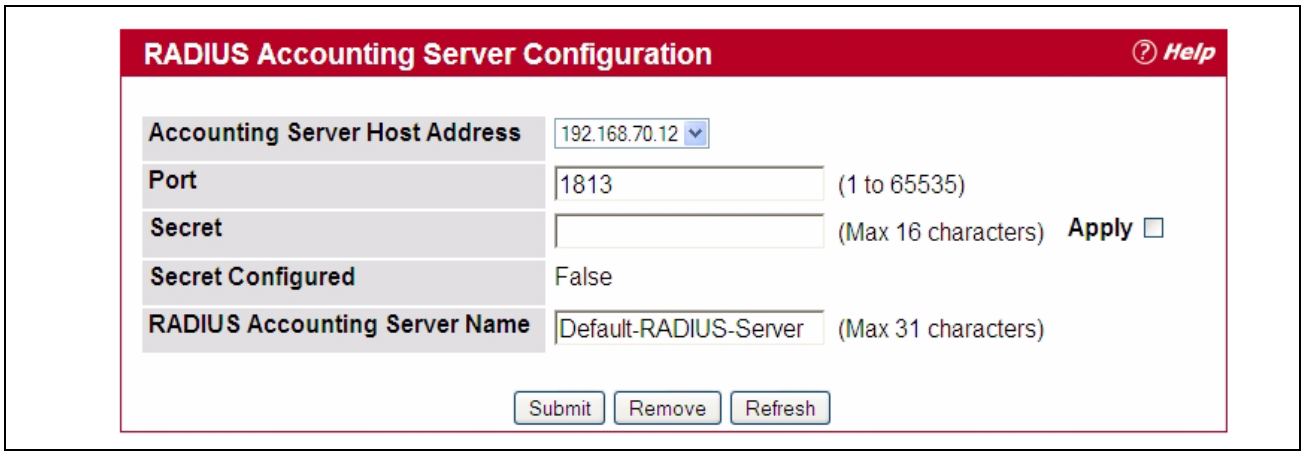

Figure 5-17: RADIUS Accounting Server Configuration—Server Added

Table 5-16: RADIUS Accounting Server Configuration Fields

Field	Description
Accounting Server Host Address	Use the drop-down menu to select the IP address of the accounting server to view or configure. Select Add to configure additional RADIUS servers.
Port	Identifies the authentication port the server uses to verify the RADIUS accounting server authentication. The port is a UDP port, and the valid range is 1-65535. The default port for RADIUS accounting is 1813.
Secret	Specifies the shared secret to use with the specified accounting server. This field is only displayed if you are logged into the switch with READ-WRITE access.
Apply	The Secret will only be applied if this box is checked. If the box is not checked, anything entered in the Secret field will have no affect and will not be retained. This field is only displayed if you are logged into the switch with READWRITE access.
Secret Configured	Indicates whether the shared secret for this server has been configured.
RADIUS Accounting Server Name	Enter the name of the RADIUS accounting server. The name can contain up to 32 alphanumeric characters. Spaces, hyphens, and underscores are also permitted. If you do not assign a name, the server is assigned the default name Default-RADIUS-Server. You can use the same name for multiple RADIUS accounting servers. RADIUS clients can use accounting servers with the same name as back-ups for each other.

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.

To delete a configured RADIUS accounting server, select the IP address of the server from the **RADIUS Server IP Address** drop-down menu, and then click **Remove**.

- Click **Refresh** to update the page with the most current information.

5.2.3.1 Named Accounting Server Status

The RADIUS Named Accounting Server Status page shows summary information about the accounting servers configured on the system.

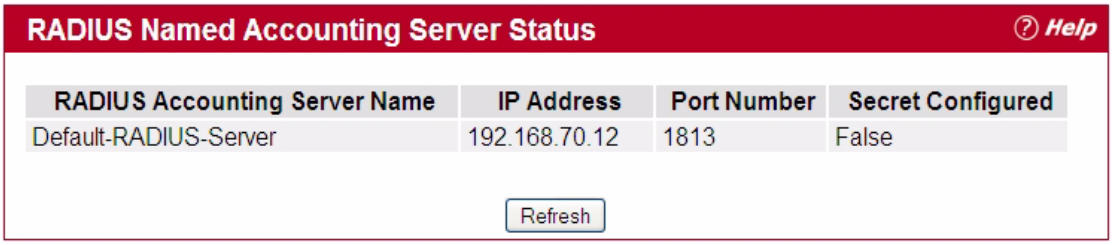

Figure 5-18: RADIUS Server Configuration—Server Added

Table 5-17: Named Accounting Server Fields

Field	Description
RADIUS Accounting Server Name	Shows the RADIUS accounting server name. Multiple RADIUS accounting servers can have the same name. In this case, RADIUS clients can use RADIUS servers with the same name as back-ups for each other.
P Address	Shows the IP address of the RADIUS server.
Port Number	Identifies the authentication port the server uses to verify the RADIUS server authentication. The port is a UDP port.
Secret Configured	Indicates whether the shared secret for this server has been configured.

Click **Refresh** to update the page with the most current information.

5.2.4 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 Server Statistics** in the navigation menu.

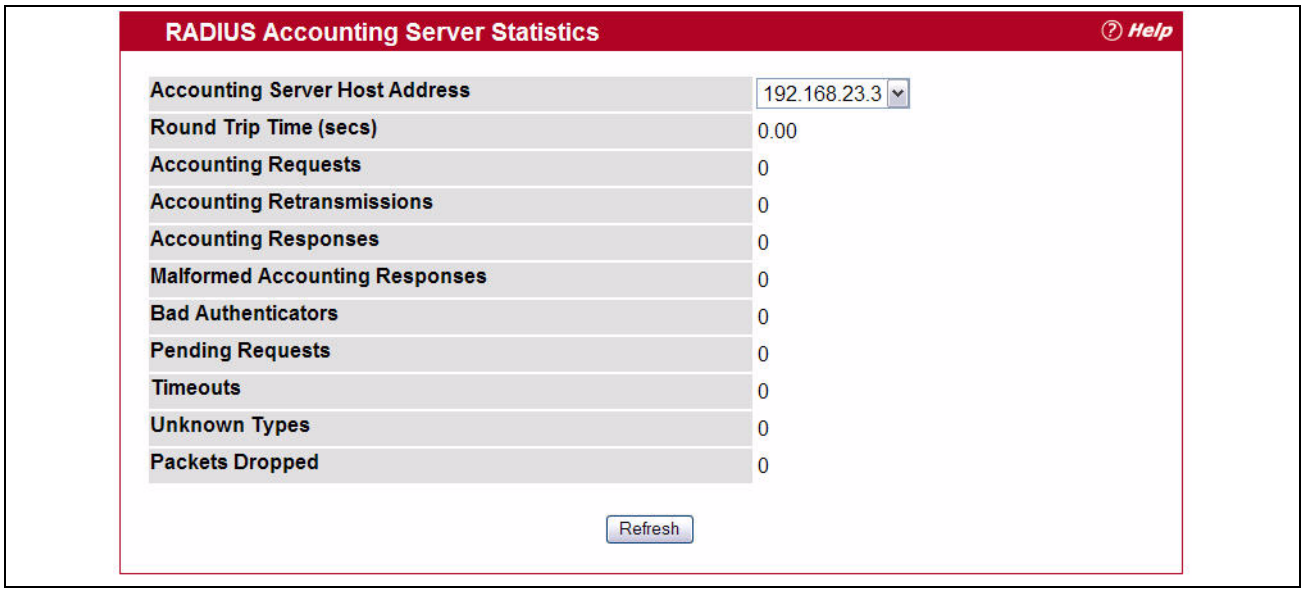

Figure 5-19: RADIUS Accounting Server Statistics

Table 5-18: RADIUS Accounting Server Fields

Field	Description
Accounting Server Host Address	Use the drop-down menu to select the IP address of the RADIUS accounting server for which to display statistics.
Round Trip Time (secs)	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.
Accounting Retransmissions	The number of RADIUS Accounting-Request packets retransmitted to this server.
Accounting Responses	Displays the number of RADIUS packets received on the accounting port from this server.
Malformed Access Responses	Displays the number of malformed RADIUS Accounting-Response packets received from this server. Malformed packets include packets with an invalid length. Bad authenticators and unknown types are not included as malformed accounting responses.
Bad Authenticators	Displays the number of RADIUS Accounting-Response packets that contained invalid authenticators received from this accounting server.
Pending Requests	The number of RADIUS Accounting-Request packets destined for this server that have not yet timed out or received a response.
Timeouts	The number of accounting timeouts to this server.
Unknown Types	The number of RADIUS packets of unknown type which were received from this server on the accounting port.
Packets Dropped	The number of RADIUS packets received from this server on the accounting port and dropped for some other reason.

5.2.5 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.

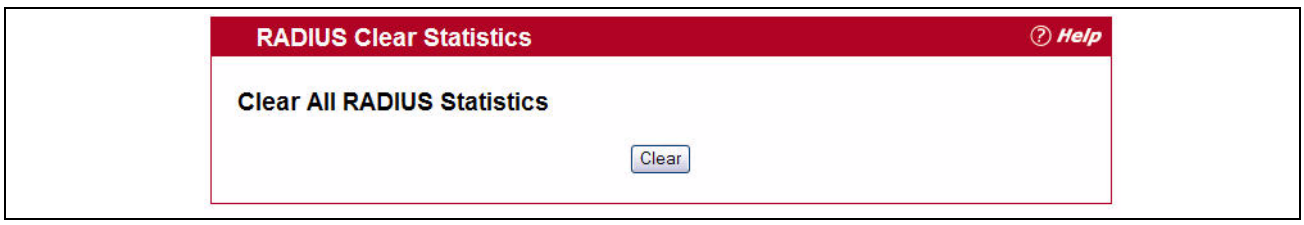

Figure 5-20: RADIUS Clear Statistics

To clear all statistics for the RADIUS authentication and accounting server, click **Clear**.

5.3 TACACS+ Settings

FASTPATH software provides Terminal Access Controller Access Control System (TACACS+) client support. TACACS+ provides centralized security for validation of users accessing the device.

TACACS+ provides a centralized user management system, while still retaining consistency with RADIUS and other authentication processes. TACACS+ provides the following services:

- **Authentication**: Provides authentication during login and via user names and user-defined passwords.
- **Authorization**: Performed at login. Once the authentication session is completed, an authorization session starts using the authenticated user name. The TACACS+ server checks the user privileges.

The TACACS+ protocol ensures network security through encrypted protocol exchanges between the device and TACACS+ server.

The TACACS+ folder contains links to the following web pages:

- TACACS+ Configuration
- TACACS+ Server Configuration

5.3.1 TACACS+ Configuration

The TACACS+ Configuration page contains the TACACS+ settings for communication between the switch and the TACACS+ server you configure. the inband management port.

To display the TACACS+ Configuration page, click **Security** > **TACACS+ > Configuration** in the navigation menu.

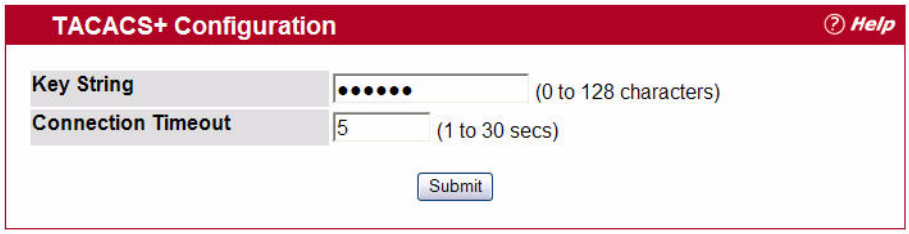

Figure 5-21: TACACS+ Configuration

Table 5-19: TACACS+ Configuration Fields

Field	Description
Key String	Specifies the authentication and encryption key for TACACS+ communications between the device and the TACACS+ server. The valid range is 0-128 characters. The key must match the key configured on the TACACS+ server.
Connection Timeout	The maximum number of seconds allowed to establish a TCP connection between the device and the TACACS+ server.

If you make any changes to the page, click **Submit** to apply the new settings to the system.

5.3.2 TACACS+ Server Configuration

Use the TACACS+ Server Configuration page to configure up to five TACACS+ servers with which the switch can communicate.

To display the TACACS+ Server Configuration page, click **Security > TACACS+ > Server Configuration** in the navigation menu.

Figure 5-22 shows the RADIUS Accounting Server Configuration page when no RADIUS servers are configured or when you select Add from the **Accounting Server IP Address** field.

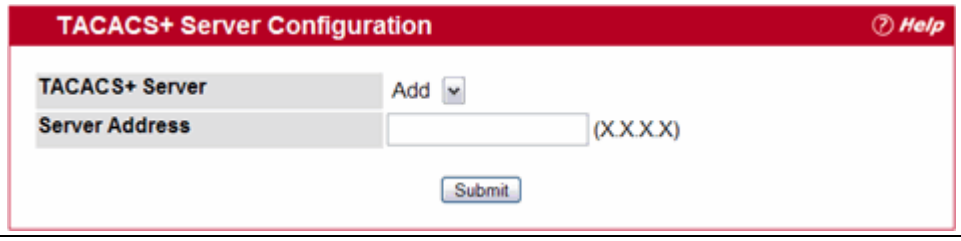

Figure 5-22: TACACS+ Configuration—No Server

After you add one or more TACACS+ servers, additional fields appear on the RADIUS Accounting Server Configuration page, as Figure 5-23 shows.

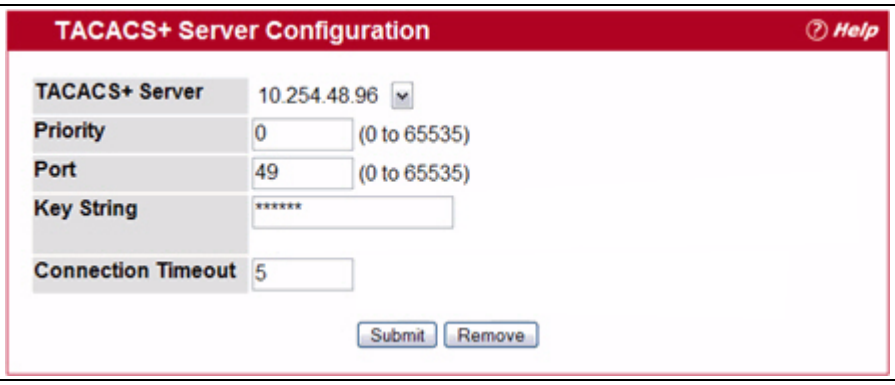

Figure 5-23: TACACS+ Configuration—Server Added

Table 5-20: TACACS+ Configuration Fields

Field	Description
TACACS+ Server	Use the drop-down menu to select the IP address of the TACACS+ server to view or configure. If fewer than five RADIUS servers are configured on the system, the Add option is also available. Select Add to configure additional RADIUS servers.

Table 5-20: TACACS+ Configuration Fields (Continued)

Field	Description
IP Address	Enter the IP address of the RADIUS accounting server to add. This field is only available when Add is selected in the **RADIUS Server IP Address** field.
Port	The authentication port number through which the TACACS+ session occurs. The default is port 49, and the range is 0-65535.
Key String	Defines the authentication and encryption key for TACACS+ communications between the device and the TACACS+ server. This key must match the encryption used on the TACACS+ server. The valid range is 0-128 characters.
Connection Timeout	The amount of time that passes before the connection between the device and the TACACS+ server times out. The field range is from 1 to 30 seconds.

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

To delete a configured TACACS+ server, select the IP address of the server from the **RADIUS Server IP Address** drop-down menu, and then click **Remove**.

5.4 Secure HTTP

Secure HTTP enables the transmission of HTTP over an encrypted Secure Sockets Layer (SSL) or Transport Layer Security (TLS) connection. When you manage the switch by using a Web interface, secure HTTP can help ensure that communication between the management system and the switch is protected from eavesdroppers and man-in-the-middle attacks.

5.4.1 Secure HTTP Configuration

Use the Secure HTTP Configuration page to configure the settings for HTTPS communication between the management station and the switch.

To display the Secure HTTP Configuration page, click **Security > Secure HTTP > Configuration** in the navigation menu.

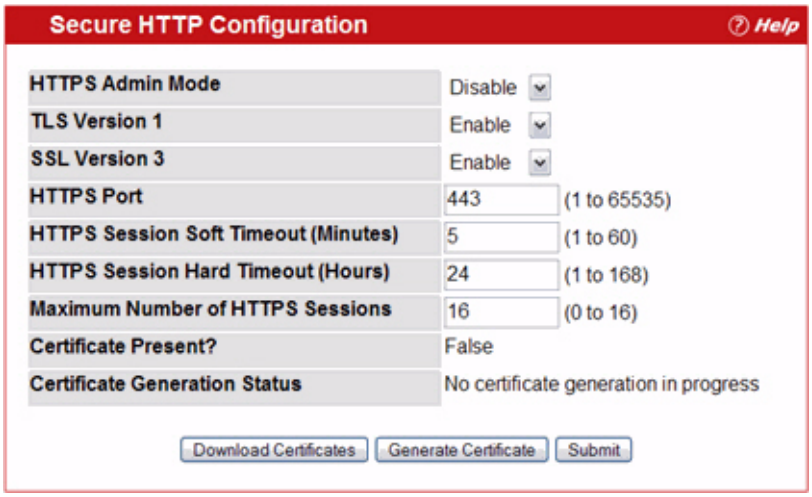

Figure 5-24: Secure HTTP Configuration

Table 5-21: Secure HTTP Configuration Fields

Field	Description
Admin Mode	Enables or Disables the Administrative Mode of Secure HTTP. The currently configured value is shown when the web page is displayed. The default value is Disable. You can only download SSL certificates when the HTTPS Admin mode is disabled.
TLS Version 1	Enables or Disables Transport Layer Security Version 1.0. The currently configured value is shown when the web page is displayed. The default value is Enable.
SSL Version 3	Enables or Disables Secure Sockets Layer Version 3.0. The currently configured value is shown when the web page is displayed. The default value is Enable.
HTTPS Port Number	Sets the HTTPS Port Number. The value must be in the range of 1 to 65535. Port 443 is the default value. The currently configured value is shown when the web page is displayed.
HTTPS Session Soft Timeout	Sets the inactivity timeout for HTTPS sessions. The value must be in the range of (1 to 60) minutes. The default value is 5 minutes. The currently configured value is shown when the web page is displayed.
HTTPS Session Hard Timeout	Sets the hard timeout for HTTPS sessions. This timeout is unaffected by the activity level of the session. The value must be in the range of (1 to 168) hours. The default value is 24 hours. The currently configured value is shown when the web page is displayed.
Maximum Number of HTTPS Sessions	Sets the maximum allowable number of HTTPS sessions. The value must be in the range of (0 to 16). The default value is 16. The currently configured value is shown when the web page is displayed.

For the Web server on the switch to accept HTTPS connections from a management station, the Web server needs a public key certificate. The switch can generate its own certificates, or you can generate these externally (i.e., off-line) and download them to the switch.

Generating Certificates

To have the switch generate the certificates:

1. Click **Generate Certificates**.

   The page refreshes with the message “Certificate generation in progress”.

2. Click **Submit** to complete the process.

   The page refreshes with the message “No certificate generation in progress” and the Certificate Present field displays as “True”.

Downloading SSL Certificates

Before you download a file to the switch, the following conditions must be true:

- The file to download from the TFTP server is on the server in the appropriate directory.
- The file is in the correct format.
- The switch has a path to the TFTP server.

Use the following procedures to download an SSL certificate.

1. Click the **Download Certificates** button at the bottom of the page.

Note...

The **Download Certificates** button is only available if the HTTPS admin mode is disabled. If the mode is enabled, disable it and click Submit. When the page refreshes, the **Download Certificates** button appears.

The Download Certificates button links to the File Download page, as Figure 5-25 shows.

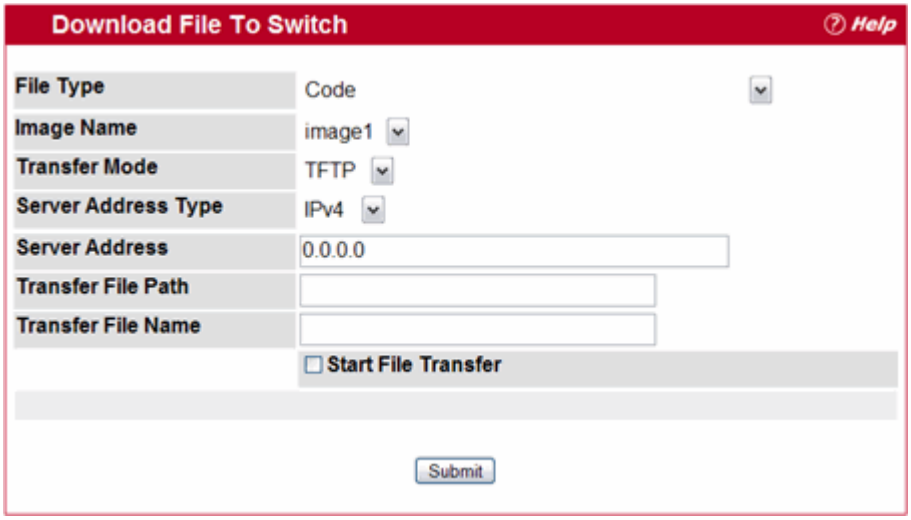

Figure 5-25: File Download

2. From the **File Type** field on the File Download page, select one of the following types of SSL files to download:
   - SSL Trusted Root Certificate PEM File: SSL Trusted Root Certificate File (PEM Encoded).
   - SSL Server Certificate PEM File: SSL Server Certificate File (PEM Encoded).
   - SSL DH Weak Encryption Parameter PEM File: SSL Diffie-Hellman Weak Encryption Parameter File (PEM Encoded).

- SSL DH Strong Encryption Parameter PEM File: SSL Diffie-Hellman Strong Encryption Parameter File (PEM Encoded).

3. Verify the IP address of the TFTP server and ensure that the software image or other file to be downloaded is available on the TFTP server.
4. Complete the **TFTP Server IP Address** and **TFTP File Name** (full path without TFTP server IP address) fields.
5. Select the **Start File Transfe**r check box, and then click **Submit**.

   After you click Submit, the screen refreshes and a “File transfer operation started” message appears. After the software is downloaded to the device, a message appears indicating that the file transfer operation completed successfully.
6. To return to the Secure HTTP Configuration page, click **Security > Secure HTTP > Configuration** in the navigation menu.
7. To enable the HTTPS admin mode, select Enable from the **HTTPS Admin Mode** field, and then click **Submit**.

5.5 Secure Shell

If you use the command-line interface (CLI) to manage the switch from a remote system, you can use Secure Shell (SSH) to establish a secure connection. SSH uses public-key cryptography to authenticate the remote computer.

5.5.1 Secure Shell Configuration

Use the Secure Shell Configuration page to configure the settings for secure command-line based communication between the management station and the switch.

To display the Secure Shell Configuration page, click **Security > Secure Shell > Configuration** in the navigation menu.

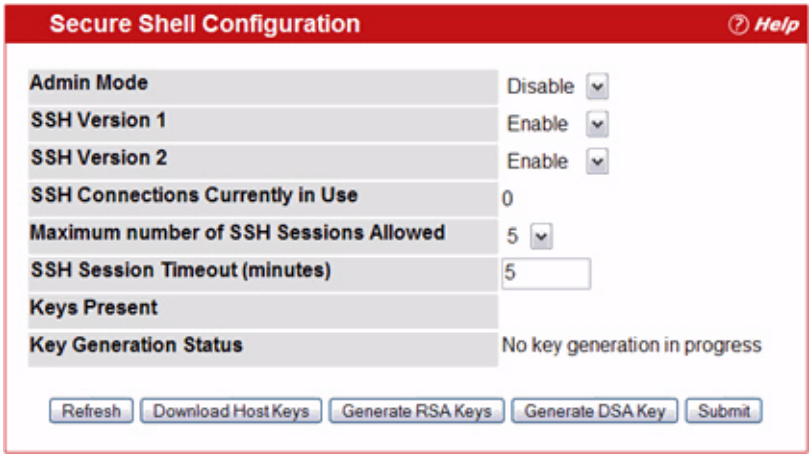

Figure 5-26: Secure Shell Configuration

Table 5-22: Secure Shell Configuration Fields

Field	Description
Admin Mode	This select field is used to Enable or Disable the administrative mode of SSH. The currently configured value is shown when the web page is displayed. Setting this value to disable shuts down the SSH port. If the admin mode is set to disable, then all existing SSH connections remain connected until timed-out or logged out, but new SSH connections cannot be established. The default value is Disable.
SSH Version 1	This select field is used to Enable or Disable Protocol Level 1 for SSH. The currently configured value is shown when the web page is displayed. The default value is Enable.
SSH Version 2	This select field is used to Enable or Disable Protocol Level 2 for SSH. The currently configured value is shown when the web page is displayed. The default value is Enable.
SSH Connections in Use	Displays the number of SSH connections currently in use in the system.
Maximum Number of SSH Sessions Allowed	This select field is used to configure the maximum number of inbound SSH sessions allowed on the switch. The currently configured value is shown when the web page is displayed. The range of acceptable values for this field is (0-5).
SSH Session Timeout (Minutes)	This text field is used to configure the inactivity timeout value for incoming SSH sessions to the switch. The acceptable range for this value is (1-160) minutes.

5.5.1.1 Downloading SSH Host Keys

For the switch to accept SSH connections from a management station, the switch needs SSH host keys or certificates. The switch can generate its own keys or certificates, or you can generate these externally (i.e., offline) and download them to the switch.

To download an SSH host key from a TFTP server to the switch, use the instructions in Downloading SSL Certificates361. However, from the File Type field on the File Download page, select one of the following key file types to download:

- **SSH-1 RSA Key File:** SSH-1 Rivest-Shamir-Adleman (RSA) Key File.
- **SSH-2 RSA Key PEM File:** SSH-2 Rivest-Shamir-Adleman (RSA) Key File (PEM Encoded).
- **SSH-2 DSA Key PEM File:** SSH-2 Digital Signature Algorithm (DSA) Key File (PEM Encoded).

RADIUS Named Server Status Help

Current	RADIUS Server IP Address	RADIUS Server Name	Port Number	Server Type	Secret Configured	Message Authenticator
True	10.27.65.66	Default-RADIUS-Server	1812	Secondary	No	Enable

Refresh

6 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 tree menu. This section contains the following subsections:

- Configuring Access Control Lists
- Configuring Differentiated Services
- Configuring Class of Service
- Configuring Auto VoIPConfiguring iSCSI Optimization

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.

Note...

Some of the features described in this section may not be supported in FATPATH software releases for particular hardware platforms.

6.1 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. FASTPATH software supports IPv4, IPv6, and MAC ACLs. The total number of MAC and IP ACLs supported by FASTPATH software is platform-specific.

The Access Control Lists folder contains links to the following folders and web pages:

- IP Access Control Lists
- IPv6 Access Control Lists
- MAC Access Control Lists
- ACL Interface Configuration
- VLAN ACL Configuration
- ACL Interface/VLAN Summary

You first create an IPv4-based, IPv6-based, or MAC-based rule and assign a unique ACL ID. Then, you define the rules, which can identify protocols, source and destination IP and MAC addresses, and other packet-matching criteria. Finally, you use the ID number to assign the ACL to a port or to a VLAN interface.

6.1.1 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 the following web pages that allow you to configure and view IP ACLs:

- IP ACL Configuration
- IP ACL Summary
- IP ACL Rule Configuration

First, you use the IP ACL Configuration page to define the IP ACL type and assign an ID to it. Then, you use the IP ACL Rule Configuration page to create rules for the ACL. Finally, you use the ACL Interface Configuration and/or ACL Interface/VLAN Summary pages to assign the ACL by its ID number to a port or VLAN. You can use the IP ACL Summary page to view a the configurations.

6.1.1.1 IP ACL Configuration

Use the IP ACL Configuration page to add or remove IP-based ACLs. On this menu the interfaces to which an IP ACL applies must be specified, as well as whether it applies to inbound or outbound traffic. Rules for the IP ACL are specified/created using the IP ACL Rule Configurationpage.

To display the IP ACL Configuration page, click **QoS > Access Control Lists > IP Access Control Lists > Configuration** in the navigation menu.

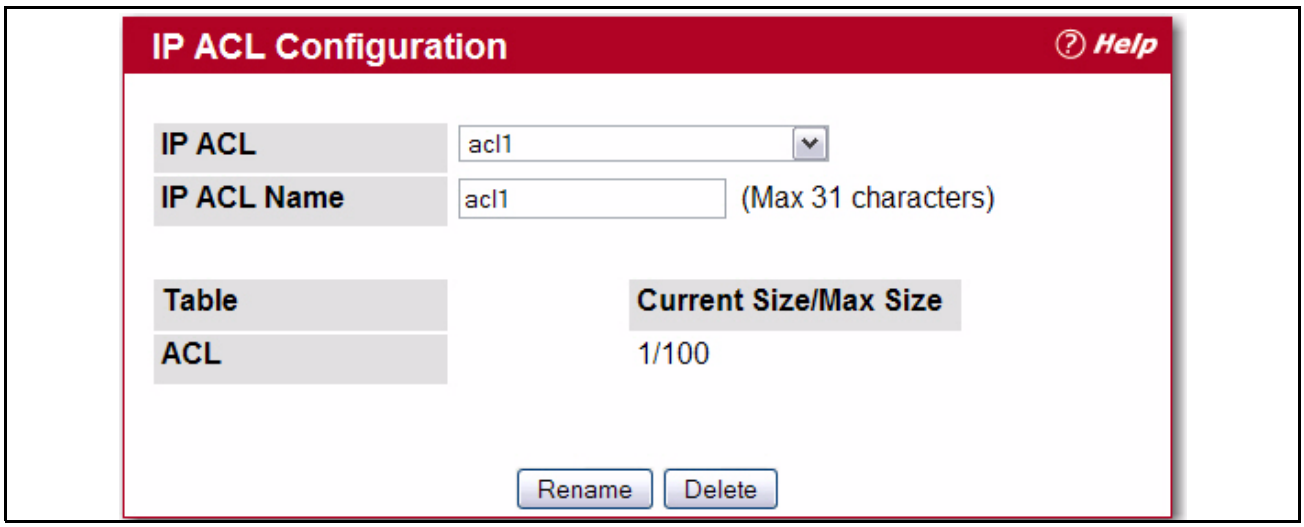

Figure 6-1: IP ACL Configuration

Table 6-1: IP ACL Configuration Fields

Field	Description
IP ACL	Select a type of ACL to create, or select an existing ACL to delete from the dropdown menu. You can create the following types of IP ACLs: • **Standard IP ACL**: Allows you to permit or deny traffic from a source IP address. • **Extended IP ACL**: Allows you to permit or deny specific types of layer 3 or layer 4 traffic from a source IP address to a destination IP address. This type of ACL provides more granularity and filtering capabilities than the standard IP ACL. • **Named IP ACL**: Allows you to create an Extended IP ACL that is identified by a name rather than a number. These ACLs have the same capabilities as Extended IP ACLs with respect to match criteria and actions supported.
IP ACL ID	Enter an ID number for the ACL to configure. This field appears if you select Create Standard IP ACL or Create Extended IP ACL from the **IP ACL** dropdown menu. For a standard IP ACL, the acceptable ID values are 1-99. For an extended IP ACL, the acceptable ID values are 101-199.
IP ACL Name	This field appears if you select Create New Named IP ACL from the **IP ACL** dropdown menu. Specify an IP ACL Name string which includes only alphanumeric characters. The name must start with an alphabetic character. This field will display the name of the currently selected IP ACL if the ACL has already been created.

The ACL Table at the bottom of the page shows the current size of the ACL table versus the maximum size of the ACL table. The current size is equal to the number of configured IPv4 and IPv6 ACLs plus the number of configured MAC ACLs. The maximum size is 100.

- To add an IP ACL, select the type of ACL to add from the **IP ACL** dropdown menu, enter an ACL ID in the appropriate field, and then click **Submit**.
- To delete an IP ACL, select the ACL ID from the **IP ACL** dropdown menu, and then click **Delete**. The **Delete** button only appears if a configured IP ACL is selected.

6.1.1.2 IP ACL Summary

Use the IP ACL Summary page to view all IP ACLs and their related data.

To display the IP ACL **Summary** page, click **QoS > Access Control Lists > IP Access Control Lists > Summary** in the navigation menu.

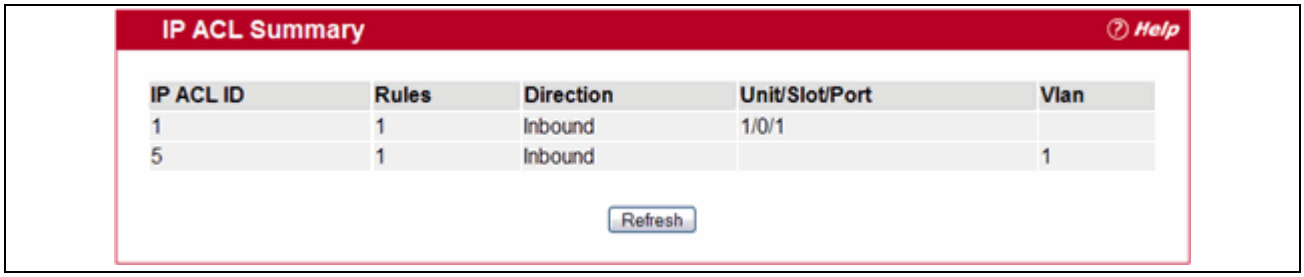

Figure 6-2: IP ACL Summary

Table 6-2: IP ACL Summary Fields

Field	Description
IP ACL ID/Name	Shows the user-configured ID number/name associated with the ACL. The range for a standard IP ACL is 1-99. For an extended IP ACL, the ID range is 101-199. ACLs identified by a name are extended ACLs.
Rules	Shows the number of rules currently configured for the IP ACL.
Direction	Shows the direction of packet traffic affected by the IP ACL. Direction can only be **Inbound** or **Outbound**.
Slot/Port	Shows the interfaces to which the IP ACL applies. To apply an ACL to an interface, see 6.1.4ACL Interface Configuration393.
VLAN	VLAN(s) to which the IP ACL applies.

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

6.1.1.3 IP ACL Rule Configuration

Use the **IP ACL Rule Configuration** page to define rules for IP-based ACLs created using the IP Access Control List Configuration page. The access list definition includes rules that specify whether traffic matching the criteria is forwarded normally or discarded. Additionally, you can specify to assign traffic to a particular queue, filter on some traffic, change VLAN tag, shut down a port, and/or redirect the traffic to a particular port.

Note...

There is an implicit "deny all" rule at the end of an ACL list. This means that if an ACL is applied to a packet and if none of the explicit rules match, then the final implicit "deny all" rule applies and the packet is dropped.

To display the IP ACL **Rule Configuration** page, click **QoS > Access Control Lists > IP Access Control Lists > Rule Configuration** in the navigation menu.

The fields available on the page depend on whether you select a standard, extended, or named IP ACL from the IP ACL field, whether the rule action is permit or deny, and whether you select Create Rule or an existing rule from the Rule field.

Figure 6-3 shows the fields available when Create Rule is selected in the **Rule** field.

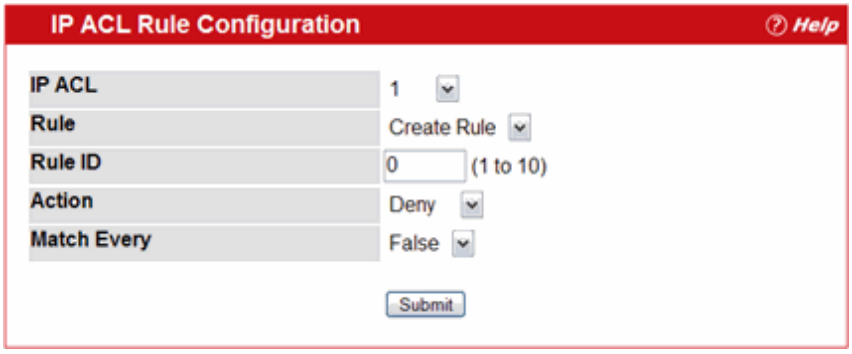

Figure 6-3: IP ACL Rule Configuration (Create Rule)

Figure 6-4 shows the fields available when you configure a rule for a standard IP ACL.

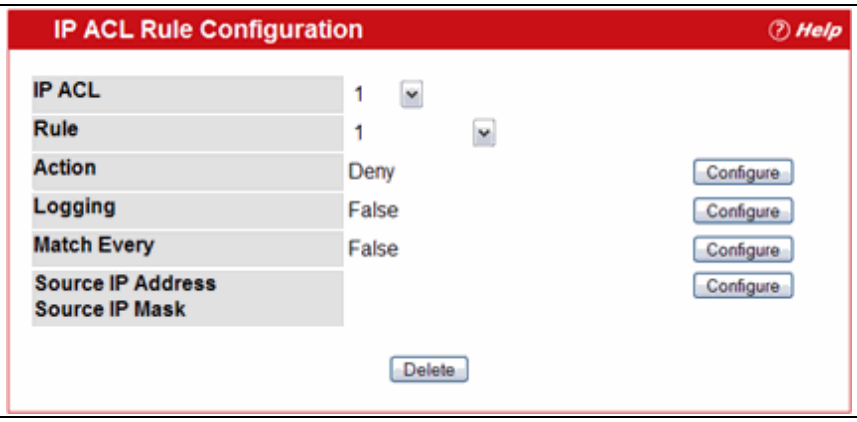

Figure 6-4: IP ACL Rule Configuration (Standard ACL)

Figure 6-5 shows the fields available when you create a rule for an extended IP ACL.

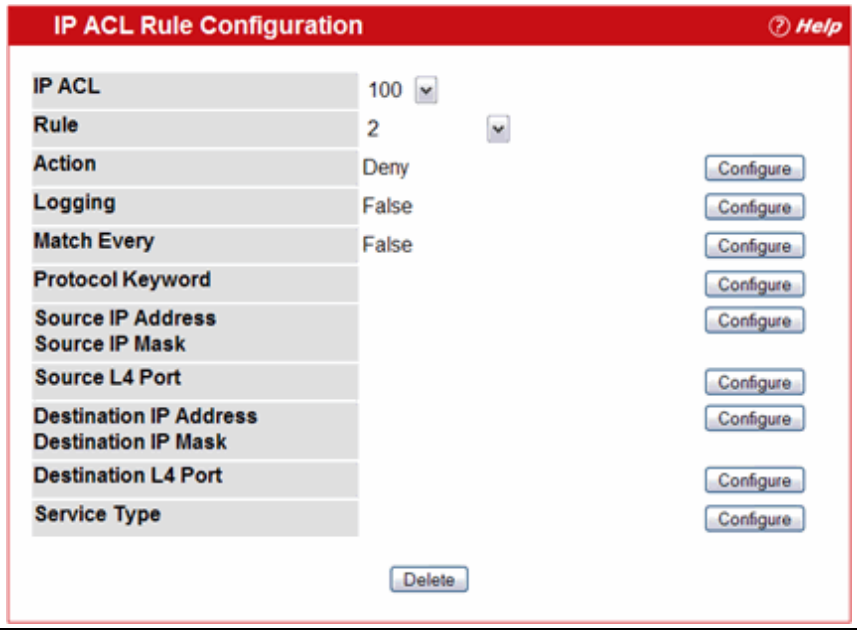

Figure 6-5: IP ACL Rule Configuration (Extended ACL Rule)

Figure 6-6 shows the fields available when you create a rule for a Named IP ACL. A Named IP ACL is identical to an Extended IP ACL in every characteristic and capability except that the Named IP ACL is identified by name rather than number.

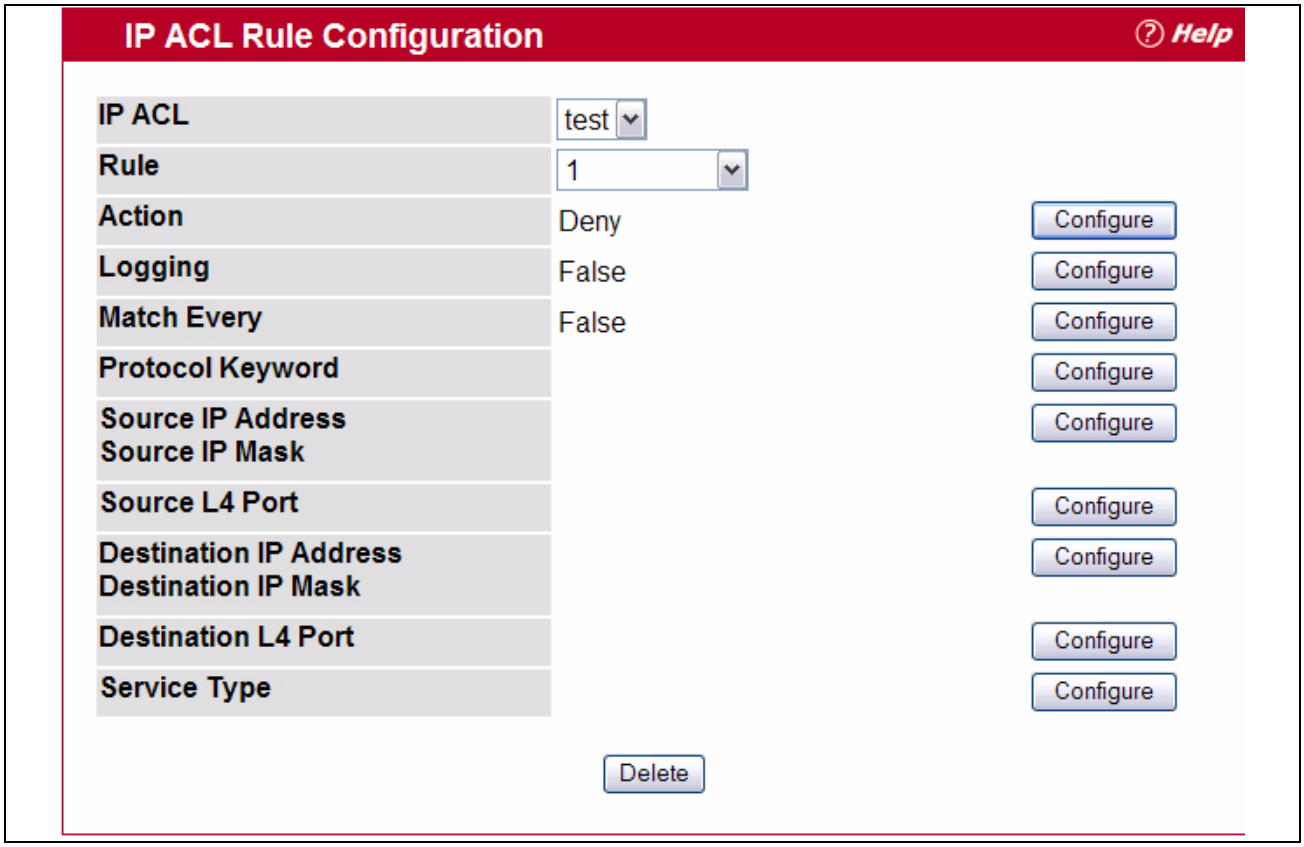

Figure 6-6: IP ACL Rule Configuration (Named ACL Rule)

Table 6-3 shows all possible fields on the IP ACL Rule Configuration page. The actual fields available on the page depend on what type of rule you configure, whether you create a new rule or modify an existing rule, and whether the rule action is Permit or Deny.

Table 6-3: IP ACL Rule Configuration Fields

Field	Description
IP ACL	The menu contains the existing IP ACLs configured on the page. To set up a new IP ACL, see 6.1.1IP Access Control Lists376.
Rule	Select an existing Rule ID to modify or select Create Rule to configure a new ACL Rule. New rules cannot be created if the maximum number of rules has been reached. for each rule, a packet must match all the specified criteria in order to be true against that rule and for the specified rule action (Permit/Deny) to take place.
Rule ID	This field is only available if you select Create Rule from the Rule field. Enter a new Rule ID which is a whole number in the range of 1 to 28 that will be used to identify the rule. After you click **Submit**, the new ID is created and you can configure the rule settings. The number of rules you can create in an ACL is platform dependent.
Action	Selects the ACL forwarding action. Click **Configure** to change the action. Select the desired action from the dropdown menu, and then click **Submit** or **Cancel** to return to the Rule Configuration page. Possible values are; • **Permit.** Forwards packets which meet the ACL criteria. • **Deny.** Drops packets which meet the ACL criteria.

Table 6-3: IP ACL Rule Configuration Fields (Continued)

Field	Description
Logging	This field is only visible for a Deny Action. When set to True, 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.
Assign Queue ID	This field is only visible when the Action is Permit. Use this field to specify the hardware egress queue identifier used to handle all packets matching this AP ACL Rule. Click **Configure**, and then enter an identifying queue number in the appropriate field. The number of queues available to select from is dependent upon platform. Click **Submit** or **Cancel** to return to the Rule Configuration page.
Redirect Interface	This field is only visible when the Action is Permit. Use this field to specify the specific egress interface where the matching traffic stream is forced, bypassing any forwarding decision normally performed by the device. This field cannot be set if a Mirror Interface is already configured for the ACL rule. Click **Configure**, and then select an interface from the dropdown list. Packets that meet the rule are redirected to the interface you select. Click **Submit** or **Cancel** to return to the Rule Configuration page.
Mirror Interface	This field is only visible when the Action is Permit. Use this field to specify the specific egress interface where the matching traffic stream is copied in addition to being forwarded normally by the device. Click **Configure**, and then select an interface from the dropdown list. Packets that meet the rule are mirrored on the interface you select. Click **Submit** or **Cancel** to return to the Rule Configuration page.
Match Every	Requires a packet to match the criteria of this ACL. Click **Configure**, and then select True or False from the dropdown list. Then click **Submit** or **Cancel** to return to the Rule Configuration page. True signifies that all packets will match the selected IP ACL and Rule and will be either permitted or denied. Match Every is exclusive to the other filtering rules, so if Match Every is True, the other rules on the screen do not appear. To configure specific match criteria for the rule, remove the rule and re-create it, or reconfigure ‘Match Every’ to ‘False’ for the other match criteria to be visible.
Protocol Keyword	Specify that a packet’s IP protocol is a match condition for the selected IP ACL rule. The possible values are ICMP, IGMP, IP, TCP, and UDP. Either the ‘Protocol Keyword’ field or the ‘Protocol Number’ field can be used to specify an IP protocol value as a match criteria. Click **Configure**, and then select the protocol keyword from the dropdown list. Click **Submit** or **Cancel** to return to the Rule Configuration page.
Protocol Number	Specify that a packet’s IP protocol is a match condition for the selected IP ACL rule and identify the protocol by number. The protocol number is a standard value assigned by IANA and is interpreted as a integer from 1 to 255. Either the ‘Protocol Number’ field or the ‘Protocol Keyword’ field can be used to specify an IP protocol value as a match criteria.
Source IP Address	Requires a packet’s source port IP address to match the address listed here. Click **Configure**, and then enter an IP Address in the appropriate field using dotted-decimal notation. The address you enter is compared to a packet's source IP Address. You also configure the Source IP Mask on the page.
Source IP Mask	Specifies the source IP address wildcard mask. Wild card masks determines which bits 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. After you enter the desired information for the Source IP Address and Source IP Mask, click **Submit** or **Cancel** to return to the Rule Configuration page.

Table 6-3: IP ACL Rule Configuration Fields (Continued)

Field	Description
Source L4 Port	Requires a packet's TCP/UDP source port to match the port listed here. Click **Configure** access the configuration page, then complete one of the following fields: • **Source L4 Keyword**: Select the desired L4 keyword from a list of source ports on which the rule can be based. If you select a keyword other than Other, the screen refreshes and the **Source L4 Port Number** field disappears. • **Source L4 Port Number**: If the source L4 keyword is Other, enter a user-defined Port ID by which packets are matched to the rule.
Destination IP Address	Requires a packet's destination port IP address to match the address listed here. Click **Configure**, and then enter an IP Address in the appropriate field using dotted-decimal notation. The address you enter is compared to a packet's destination IP Address. You also configure the Destination IP Mask on the page.
Destination IP Mask	Specify the IP mask in dotted-decimal notation to be used with the Destination IP Address value.
Destination L4 Port	Requires a packet's TCP/UDP destination port to match the port listed here. Click **Configure** access the configuration page, then complete one of the following fields: • **Destination L4 Keyword**: Select the desired L4 keyword from a list of destination ports on which the rule can be based. If you select a keyword other than Other, the screen refreshes and the **Destination L4 Port Number** field disappears. • **Destination L4 Port Number**: If the destination L4 keyword is Other, enter a user-defined Port ID by which packets are matched to the rule.
Service Type	Select one of the following three Match conditions for the extended IP ACL rule. These are alternative ways of specifying a match condition for the same Service Type field in the IP header, however each uses a different user notation. After a selection is made, the appropriate value can be specified: • **IP DSCP**: Matches the packet DSCP value to the rule. Specify the IP DiffServ Code Point (DSCP) field. The DSCP is defined as the high-order six bits of the Service Type octet in the IP header. This is an optional configuration. Enter an integer from 0 to 63. Either the DSCP value or the IP Precedence value is used to match packets to ACLs. The IP DSCP is selected by selecting one of the DSCP keyword values from a dropdown menu. If a value is to be selected by specifying its numeric value, then select the 'Other' option in the dropdown menu and a text box will appear where you can enter the numeric value of the DSCP. • **IP Precedence**: The IP Precedence field in a packet is defined as the high-order three bits of the Service Type octet in the IP header. This is an optional configuration. Matches the packet IP Precedence value to the rule when checked. Enter the IP Precedence value, an integer from 0 to 7, to match. Either the DSCP value or the IP Precedence value is used to match packets to ACLs. • **IP TOS Bits**: The IP TOS field in a packet is defined as all eight bits of the Service Type octet in the IP header. Matches on the Type of Service bits in the IP header when checked. 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. This is an optional configuration. - **TOS Bits**: This value is a hexadecimal number from 00 to FF. Requires the bits in a packet's TOS field to match the two-digit hexadecimal number entered here. - **TOS Mask**: This value is a hexadecimal number from 00 to FF. Specifies the bit positions that are used for comparison against the IP TOS field in a packet.

Modifying an IP-based Rule

Note...

Rules can be modified only when the ACL to which they belong is not bound to an interface.

1. Open the **IP ACL Rule Configuration** page.
2. Select **the desired ACL from the IP ACL dropdown menu**.
3. Select the desired rule from the **Rule ID** dropdown menu.
4. Modify the remaining fields as needed.
5. Click **Submit**.

   The IP-based rule is modified, and the device is updated.

Adding a New Rule to an IP-based ACL

1. Open the **IP ACL Rule Configuration** page.
2. Select **the desired ACL from the IP ACL dropdown menu**.
3. Specify Create Rule for **Rule ID** and enter a new ID number.
4. Define the remaining fields as needed.
5. Click **Submit**.

   The new rule is assigned to the specified IP-based ACL.

Deleting a Rule from an IP-based ACL

1. Open the **IP ACL Rule Configuration** page.
2. Select **the desired ACL from the IP ACL dropdown menu**.
3. Select the rule to delete from the **Rule** field.
4. Click **Delete**.

   The new rule is assigned to the specified IP-based ACL.
5. Click **Refresh** to update the page with the most current information.

6.1.2 IPv6 Access Control Lists

An IPv6 ACL consists of a set of rules which are matched sequentially against a packet. When a packet meets the match criteria of a rule, the specified rule action (Permit/Deny) is taken and the additional rules are not checked for a match. On this menu the interfaces to which an IPv6 ACL applies must be specified, as well as whether it applies to inbound or outbound traffic. Rules for the IPv6 ACL are specified/created using the IPv6 ACL Rule Configuration menu.

The IP Access Control List folder contains links to the following web pages that allow you to configure and view IP ACLs:

- IPv6 ACL Configuration
- IPv6 ACL Summary
- IPv6 ACL Rule Configuration

First, you use the IPv6 ACL Configuration page to define the IP ACL type and assign an ID to it. Then, you use the IPv6 ACL Rule Configuration page to create rules for the ACL. Finally, you use the ACL Interface Configuration and/or ACL Interface/VLAN Summary pages to assign the ACL by its ID number to a port or VLAN. You can use the IPv6 ACL Summary page to view a the configurations.

6.1.2.1 IPv6 ACL Configuration

Use the IP ACL Configuration page to add or remove IP-based ACLs. To display the IP ACL Configuration page, click **QoS > Access Control Lists > IPv6 Access Control Lists > Configuration** in the navigation menu.

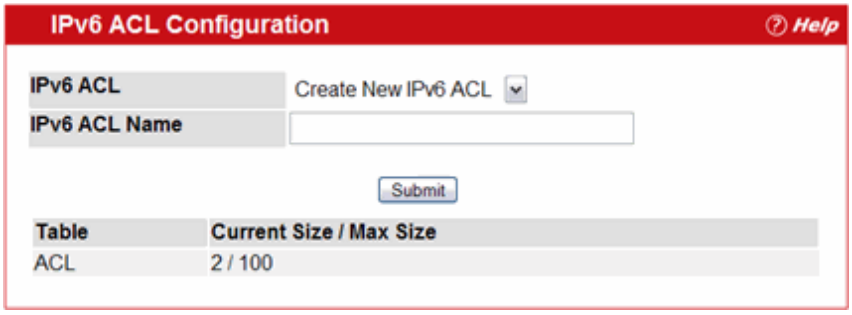

Figure 6-7: IPv6 ACL Configuration

Table 6-4: IPv6 ACL Configuration Fields

Field	Description
IPv6 ACL	Select a type of ACL to create, or select an existing ACL to delete from the dropdown menu.
IPv6 ACL Name	Specify an IPv6 ACL name string which includes alphanumeric characters only. The name must start with an alphabetic character. This field displays the name of the currently selected IPv6 ACL if the ACL has already been created.

The ACL Table at the bottom of the page shows the current size of the ACL table versus the maximum size of the ACL table. The current size is equal to the number of configured IPv4/IPv6 ACLs plus the number of configured MAC ACLs. The maximum size is 100.

- To add an IPv6 ACL, select **Create** from the IPv6 ACL list, enter an ACL name in the name, and then click **Submit**.
- To modify an IPv6 ACL name, select the name from the IPv6 ACL list, type a new name, and click **Submit**.
- To delete an IPv6 ACL, select the ACL ID from the **IP ACL** menu, and then click **Delete**. The **Delete** button appears only when a configured IP ACL is selected.

6.1.2.2 IPv6 ACL Summary

Use the IP ACL Summary page to view all IP ACLs and their related data.

To display the IP ACL **Summary** page, click **QoS > Access Control Lists > IPv6 IP Access Control Lists > Summary** in the navigation menu.

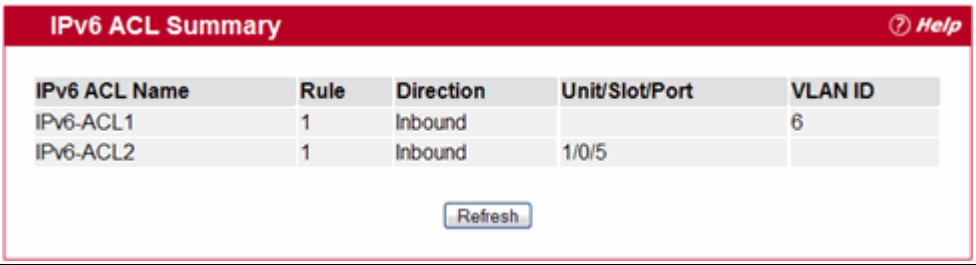

Figure 6-8: IPv6 ACL Summary

Table 6-5: IPv6 ACL Summary Fields

Field	Description
IPv6 ACL Name	Describes the number ranges for IPv4 ACL standard versus extended. The range for a standard IP ACL is 1-99. For an extended IP ACL, the ID range is 101-199.
Rules	Shows the number of rules currently configured for the IP ACL.
Direction	Shows the direction of packet traffic affected by the IP ACL, which can be Inbound or blank.
Slot/Port	Shows the interfaces to which the IP ACL applies. To apply an ACL to an interface, see 6.1.4ACL Interface Configuration393.
VLAN ID	The VLAN(s) to which the IPv6 ACL applies.

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

6.1.2.3 IPv6 ACL Rule Configuration

Use the **IPv6 ACL Rule Configuration** page to define rules for IPv6-based ACLs. The access list definition includes rules that specify whether traffic matching the criteria is forwarded normally or discarded. Additionally, you can specify to assign traffic to a particular queue, filter on some traffic, change VLAN tag, shut down a port, and/or redirect the traffic to a particular port. By default, no specific value is in effect for any of the IPv6 ACL rules.

Note...

There is an implicit "deny all" rule at the end of an ACL list. This means that if an ACL is applied to a packet and if none of the explicit rules match, then the final implicit "deny all" rule applies and the packet is dropped.

To display the IPv6 ACL **Rule Configuration** page, click **QoS > Access Control Lists > IPv6 Access Control Lists > Rule Configuration** in the navigation menu.

Figure 6-3 shows the fields available when Create Rule is selected in the **Rule** field.

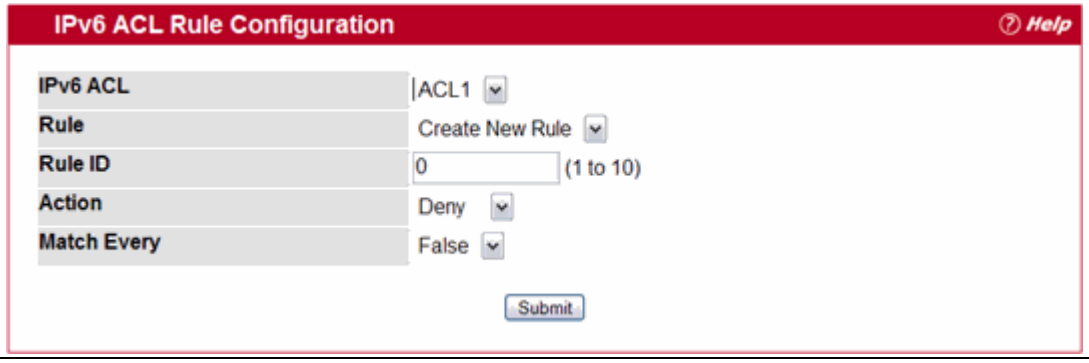

Figure 6-9: IPv6 ACL Rule Configuration (Create Rule)

Table 6-3 shows all possible fields on the IP ACL Rule Configuration page. The actual fields available on the page depend on what type of rule you configure, whether you create a new rule or modify an existing rule, and whether the rule action is Permit or Deny.

Table 6-6: IPv6 ACL Rule Configuration Fields

Field	Description
IPv6 ACL	Select the ACL you want to configure.
Rule	Select an existing Rule ID to modify or select **Create Rule** to configure a new ACL Rule. New rules cannot be created if the maximum number of rules has been reached. For each rule, a packet must match all the specified criteria in order to be true against that rule and for the specified rule action (Permit/Deny) to take place.
Rule ID	If you selected an existing rule ID in the Rule field, that ID displays here. If you are creating a new rule, then enter the next available ID number (or any other number). The number of rules you can create in an ACL is platform dependent.
Action	Specify what action should be taken if a packet matches the rule's criteria. The choices are Permit or Deny.
Logging	When set to 'True', 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 was 'hit' 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. This field is visible for a 'Deny' Action.
Match Every	Select **True** or **False** from the pulldown menu. **True** signifies that all packets will match the selected IPv6 ACL and rule and will be either permitted or denied. In this case, since all packets match the rule, the option of configuring other match criteria will not be offered. To configure specific match criteria for the rule, remove the rule and re-create it, or re-configure 'Match Every' to 'False' for the other match criteria to be visible.
Protocol	There are two ways to configure IPv6 protocol. • Specify an integer ranging from 0 to 255 after selecting protocol keyword "other". This number represents the IP protocol. • Select name of a protocol from the existing list of Internet Protocol (IP), Transmission Control Protocol (TCP), User Datagram Protocol (UDP), Internet Control Message Protocol (ICMP) and Internet Group Management Protocol (IGMP).
Additional Fields when Action = Deny	
Source Prefix/ PrefixLength	Specify IPv6 Prefix combined with IPv6 Prefix length of the network or host from which the packet is being sent. Prefix length can be in the range (0 to 128).

Table 6-6: IPv6 ACL Rule Configuration Fields (Continued)

Field	Description
Source L4 Port	Specify a packet's source layer 4 port as a match condition for the selected IPv6 ACL rule. Source port information is optional. Source port information can be specified in two ways: • Select keyword "other" from the drop down menu and specify the number of the port in the range from 0 to 65535. • Select one of the keyword from the list: DOMAIN, ECHO, FTP, FTPDATA, HTTP, SMTP, SNMP, TELNET, TFTP, and WWW. Each of these values translates into its equivalent port number, which is used as both the start and end of the port range.
Destination Prefix/ Prefix Length	Enter up to a 128-bit prefix combined with the prefix length to be compared to a packet's destination IP address as a match criteria for the selected IPv6 ACL rule. The prefix length can be in the range 0 to 128.
Destination L4 Port Number	Specify a packet's destination layer 4 port number match condition for the selected IPv6 ACL rule. This is an optional configuration.
Destination L4 Port Keyword	Specify the destination layer 4 port match conditions for the selected IPv6 ACL rule. The possible values are DOMAIN, ECHO, FTP, FTPDATA, HTTP, SMTP, SNMP, TELNET, TFTP, and WWW. Each of these values translates into its equivalent port number, which is used as both the start and end of the port range. This is an optional configuration.
Flow Label	A 20-bit number that is unique to an IPv6 packet that is used by end stations to signify QoS handling in routers. The flow label can specified within the range 0 to 1048575.
IPv6 DSCP Service	Specify the IP DiffServ Code Point (DSCP) value, which is defined as the high-order six bits of the Service Type octet in the IPv6 header. This is an optional configuration. Enter an integer from 0 to 63. The IPv6 DSCP can be selected from one of the DSCP keywords in the dropdown box. To specify a DSCP by its numeric value, select the 'Other' option in the menu and a text box displays for entering the numeric value.
Additional Fields when Action = Permit (for any field not listed below, see the Action = Deny field definitions above)	
Assign Queue ID	Specifies the hardware egress queue identifier used to handle all packets matching this IPv6 ACL rule. Valid range of Queue IDs is 0 to 6.
Mirror Interface	Specifies the egress interface where the matching traffic stream is copied, in addition to it being forwarded normally by the device. This field cannot be set if a Redirect Interface is already configured for the ACL rule.
Redirect Interface	Specifies the egress interface where the matching traffic stream is forced, bypassing any forwarding decision normally performed by the device. This field cannot be set if a Mirror Interface is already configured for the ACL rule.

- Click **Configure** to configure the corresponding match criteria for the selected rule.
- To remove the currently selected rule from the selected AC, click **Delete**.

These changes will not be retained across a power cycle unless a save configuration is performed.

6.1.3 MAC Access Control Lists

A MAC ACL consists of a set of rules which are matched sequentially against a packet. When a packet meets the match criteria of a rule, the specified rule action (Permit/Deny) is taken and the additional rules are not checked for a match. On this menu the interfaces to which an MAC ACL applies must be specified, as well as whether it applies to inbound or outbound traffic. Rules for the MAC ACL are specified/created using the MAC ACL Rule Configuration menu.

This folder links to the following pages:

- MAC ACL Configuration
- MAC ACL Summary
- MAC ACL Rule Configuration

First, you use the MAC ACL Configuration page to define the ACL type and assign an ID to it. Then, you use the MAC ACL Rule Configuration page to create rules for the ACL. Finally, you use the ACL Interface Configuration and/or ACL Interface/VLAN Summary pages to assign the ACL by its ID number to a port or VLAN. You can use the MAC ACL Summary page to view a the configurations.

6.1.3.1 MAC ACL Configuration

The MAC ACL Configuration page allows network administrators to define a MAC-based ACL. For an explanation of ACLs, see “IP Access Control Lists."

To display the MAC ACL Configuration page, click **QoS > Access Control Lists > MAC Access Control Lists > Configuration** in the navigation tree.

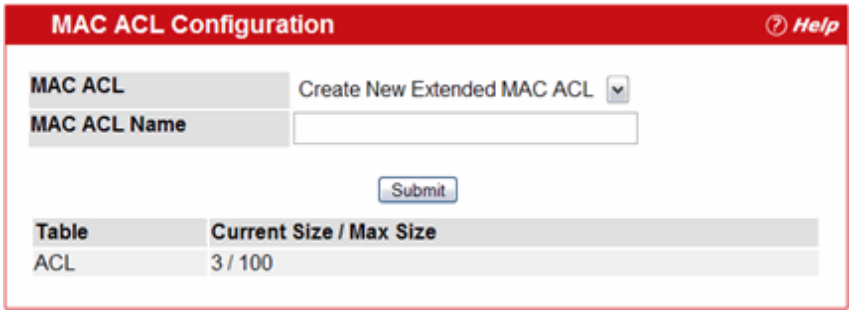

Figure 6-10: MAC ACL Configuration

Table 6-7: MAC ACL Configuration Fields

Field	Description
MAC ACL	The options in the dropdown menu allow you to create a new MAC ACL or select an existing MAC ACL that you want to rename.
MAC ACL Name	Enter a name for the MAC ACL. The name string may include alphabetic, numeric, dash, underscore, or space characters only. The name must start with an alphabetic character. This field displays the name of the currently selected MAC ACL if the ACL has already been created.

The ACL Table at the bottom of the page shows the current size of the ACL table versus the maximum size of the ACL table. The current size is equal to the number of configured IPv4/IPv6 ACLs plus the number of configured MAC ACLs. The maximum size is platform-specific.

- To add a MAC ACL, select Create New Extended MAC ACL from the **MAC ACL** dropdown menu, enter a name for the ACL in the appropriate field, and then click **Submit**.
- To rename a MAC ACL, select the ACL name from the **MAC ACL** dropdown menu. Enter a new name for the ACL in the appropriate field, and then click **Rename**. The **Rename** button only appears if a configured MAC ACL is selected.
- To delete a MAC ACL, select the ACL name from the **MAC ACL** dropdown menu, and then click **Delete**. The **Delete** button only appears if a configured MAC ACL is selected.

6.1.3.2 MAC ACL Summary

Use the MAC ACL Summary page to view all MAC ACLs and their related data.

To display the MAC ACL Summary page, click **QoS > Access Control Lists > MAC Access Control Lists > Summary** in the navigation menu.

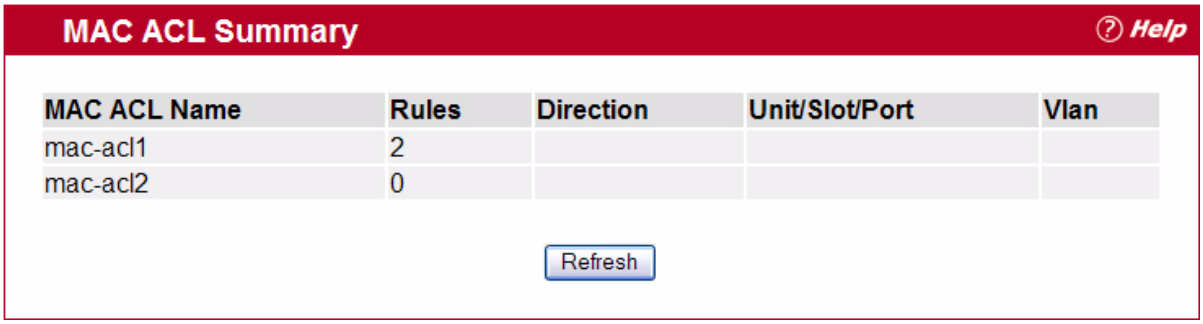

Figure 6-11: MAC ACL Summary

Table 6-8: MAC ACL Summary Fields

Field	Description
MAC ACL Name	Shows the MACL ACL Identifier.
Rules	Shows the number of rules currently configured for the MAC ACL.
Direction	Shows the direction of packet traffic affected by the MAC ACL, which can be Inbound or outbound.
Slot/Port	Shows the interfaces to which the MAC ACL applies. To apply an ACL to an interface, see 6.1.4ACL Interface Configuration393.
VLAN	VLAN(s) to which the MACL ACL applies.

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

6.1.3.3 MAC ACL Rule Configuration

Use the MAC ACL Rule Configuration page to define rules for MAC-based ACLs. The access list definition includes rules that specify whether traffic matching the criteria is forwarded normally or discarded. A default 'deny all' rule is the last rule of every list.

To display the MAC ACL Rule Configuration page, click **QoS > Access Control Lists > MAC Access Control Lists > Rule Configuration** in the navigation menu.

The fields available on the page depend on whether the rule action is permit or deny, and whether you select **Create Rule** or an existing rule from the **Rule** field.

Figure 6-12 shows the fields available when Create New Rule is selected in the **Rule** field.

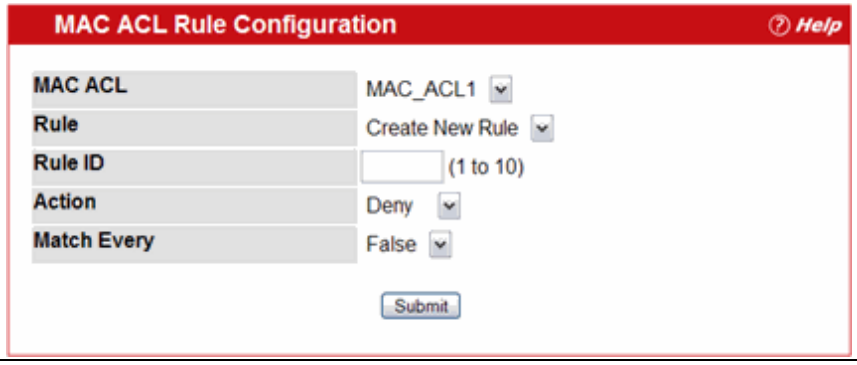

Figure 6-12: MAC ACL Rule Configuration (Create Rule)

Figure 6-13 shows the fields available when you configure a MAC ACL rule with a Deny action.

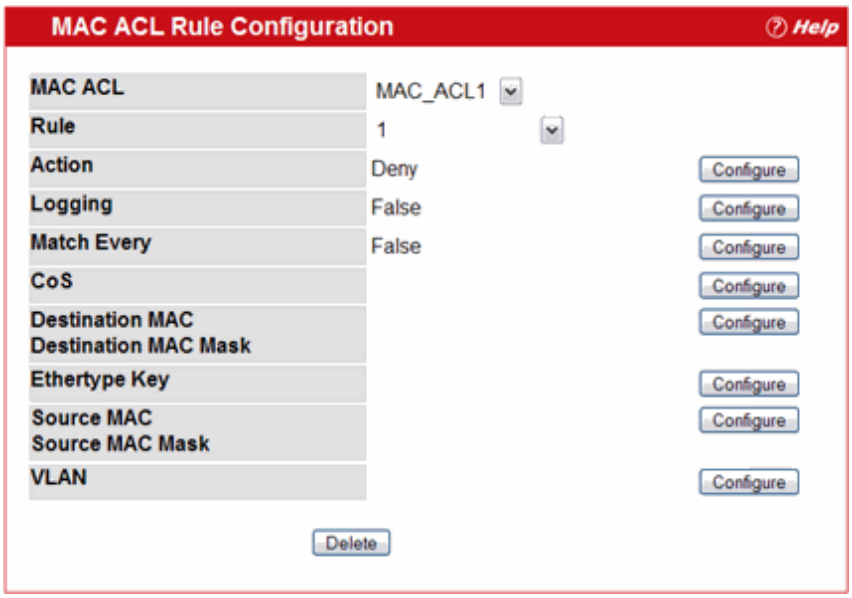

Figure 6-13: MAC ACL Rule Configuration (Deny Action)

Figure 6-14 shows the fields available when you create a rule for a MAC ACL.

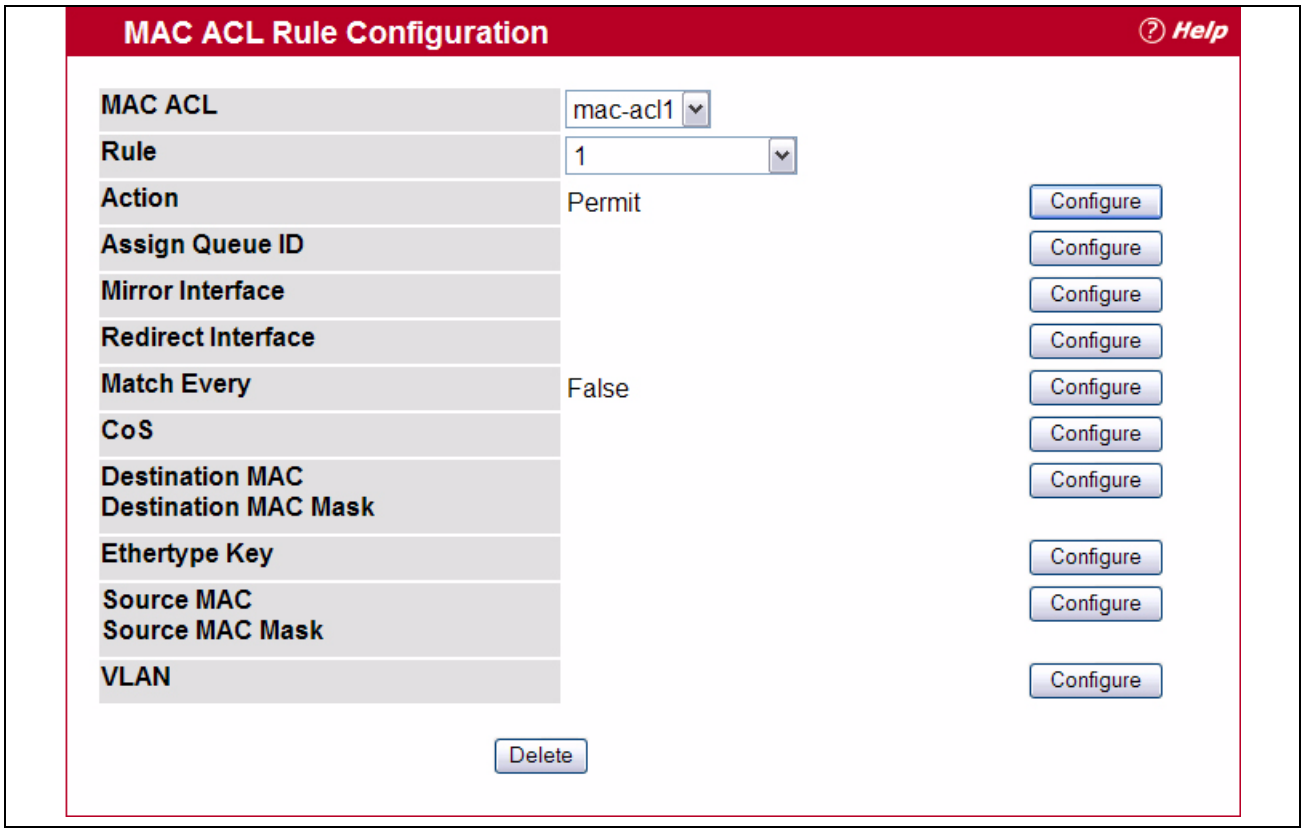

Figure 6-14: MAC ACL Rule Configuration (Permit Action)

Table 6-9 shows all possible fields on the MAC ACL Rule Configuration page. The actual fields available on the page depend on whether you create a new rule or modify an existing rule, and whether the rule action is Permit or Deny.

Table 6-9: MAC ACL Rule Configuration Fields

Field	Description
MAC ACL	Specifies an existing MAC ACL. To set up a new MAC ACL use the MAC Access Control Lists page.
Rule	Select an existing Rule ID to modify or select Create Rule to configure a new ACL Rule. Enter a whole number in the range of 1 to 28 that will be used to identify the rule. New rules cannot be created if the maximum number of rules has been reached. For each rule, a packet must match all the specified criteria in order to be true against that rule and for the specified rule action (Permit/Deny) to take place.
Rule ID	This field is only available if you select Create Rule from the Rule field. Enter a new Rule ID. After you click **Submit**, the new ID is created and you can configure the rule settings. You can create up to 10 rules for each ACL.
Action	Specify what action should be taken if a packet matches the rule's criteria: • **Permit**: Forwards packets that meet the ACL criteria. • **Deny**: Drops packets that meet the ACL criteria.

Table 6-9: MAC ACL Rule Configuration Fields (Continued)

Field	Description
Logging	This field is only visible for a Deny Action. When set to True, 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.
Assign Queue ID	This field is only visible when the Action is Permit. Specifies the hardware egress queue identifier used to handle all packets matching this ACL rule. Click Configure, and then enter an identifying number from 0 to 6 in the appropriate field. Click **Submit** or **Cancel** to return to the Rule Configuration page.
Match Every	Requires a packet to match the criteria of this ACL. Click **Configure**, and then select True or False from the dropdown list. Then click **Submit** or **Cancel** to return to the Rule Configuration page. Match Every is exclusive to the other filtering rules, so if Match Every is True, the other rules on the screen do not appear. False indicates that it is not mandatory for every packet to match the selected ACL Rule.
Mirror Interface	This field is only visible when the Action is Permit. Specifies the specific egress interface where the matching traffic stream is copied in addition to being forwarded normally by the device. This field cannot be set if a Redirect Interface is already configured for the ACL rule.
Redirect Interface	This field is only visible when the Action is Permit. Specifies the specific egress interface where the matching traffic stream is forced, bypassing any forwarding decision normally performed by the device. This field cannot be set if a Mirror Interface is already configured for the ACL rule.
CoS	Specifies the 802.1p user priority to compare against an Ethernet frame. Requires a packet's class of service (CoS) to match the CoS value listed here. Click **Configure**, and then enter a CoS value between 0 and 7 to apply this criteria. Click **Submit** or **Cancel** to return to the Rule Configuration page.
Destination MAC Address	Requires an Ethernet frame's destination port MAC address to match the address listed here. Click **Configure**, and then enter a MAC address in the appropriate field. The valid format is xx_xx_xx_xx_xx_xx. The BPDU keyword may be specified using a Destination MAC Address of 01:80:C2:xx:xx:xx. Click **Submit** or **Cancel** to return to the Rule Configuration page.
Destination MAC Mask	If desired, enter the MAC Mask associated with the Destination MAC to match. The MAC address mask specifies which bits in the destination MAC to compare against an Ethernet frame. Use F's and zeros in the MAC mask, which is in a wildcard format. An F means that the bit is not checked, and a zero in a bit position means that the data must equal the value given for that bit. For example, if the MAC address is aa_bb_cc_dd_ee_ff, and the mask is 00_00_ff_ff_ff_ff, all MAC addresses with aa_bb_xx_xx_xx_xx result in a match (where *x* is any hexadecimal number). Click **Submit** or **Cancel** to return to the Rule Configuration page.
EtherType Key	Requires a packet's EtherType to match the EtherType you select. Click **Configure**, and then select the EtherType value from the dropdown menu. If you select User Value, you can enter a custom EtherType value.
Ethertype User Value	This field only appears if you select User Value from the EtherType dropdown list. The value you enter specifies a customized Ethertype to compare against an Ethernet frame. The valid range of values is (0x0600 to 0xFFFF).
Source MAC Address	Requires a packet's source port MAC address to match the address listed here. Click **Configure**, and then enter a MAC address in the appropriate field. The valid format is xx:xx:xx:xx:xx:xx.
Source MAC Mask	If desired, enter the MAC mask for the source MAC address to match. Use F's and zeros in the MAC mask, which is in a wildcard format. An F means that the bit is not checked, and a zero in a bit position means that the data must equal the value given for that bit. The valid format is xx:xx:xx:xx:xx:xx. Click **Submit** or **Cancel** to return to the Rule Configuration page.
VLAN	Requires a packet's VLAN ID to match the ID listed here. Click **Configure**, and then enter the VLAN ID to apply this criteria. The valid range is platform specific. Either VLAN Range or VLAN can be configured. Click **Submit** or **Cancel** to return to the Rule Configuration page.

Adding a New Rule to a MAC-based ACL

Once you configure a MAC ACL, you can add rules to the ACL.

1. Open the **MAC ACL Rule Configuration** page.
2. If more than one MAC ACL is configured on the system, select the desired ACL from the MAC ACL dropdown menu.
3. From the **Rule** dropdown menu, select Create New Rule.
4. Enter a new ID number for the rule.
5. Configure the remaining rule criteria as needed.
6. Click **Submit**.

   The new rule is assigned to the specified MAC-based ACL.

Removing a Rule From a MAC-based ACL

1. From the **MAC ACL Rule Configuration** page, select an ACL from the **MAC ACL** field.
2. Select a rule from the **Rule** dropdown menu.
3. Click **Delete**.

   The rule is removed from the MAC-based ACL, and the device is updated.

6.1.4 ACL Interface Configuration

When an ACL is bound to an interface, all the rules that have been defined are applied to the selected interface. Use the ACL Interface Configuration page to assign ACLs and Interfaces and prioritize the ACLs that are bound to each interface. You can also assign ACLs to a VLAN rather than a port; see 6.1.6ACL Interface/VLAN Summary396 for information.

To display the ACL Interface Configuration page, click **QoS > Access Control Lists > Interface Configuration** in the navigation menu.

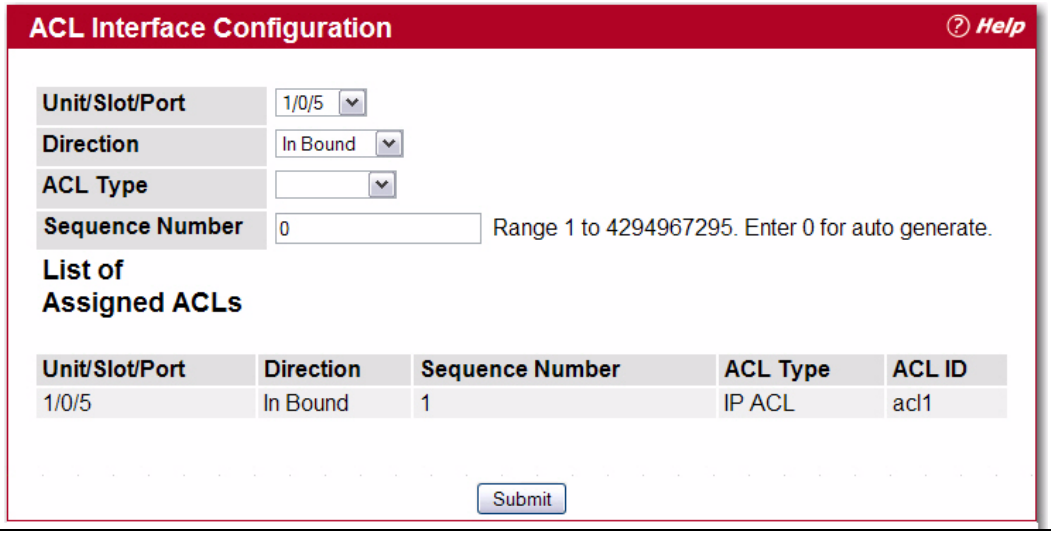

Figure 6-15: ACL Interface Configuration

If an ACL has been assigned to the interface, it displays in the table at the bottom of the page.

Table 6-10: ACL Interface Configuration Fields

Field	Description
Slot/Port	Select the interface, LAG, or VLAN routing interface from the dropdown menu.
Direction	Specifies the packet filtering direction for the ACL. The system supports Inbound and Outbound filtering. inbound filtering means the system applies the ACL rules to packets as they enter the interface.
ACL Type	Use the menu to select the ACL type to which incoming packets are matched. Packets can be matched to IP-, IPv6-, or MAC-based ACLs.
IPv4/IPv6/MAC ACL	Select the ACL of the specified type to apply to the interface from the dropdown menu.
ACL Identifier	Displays the ACL Number or Name identifying the ACL assigned to selected interface and direction.
Sequence Number	Assigns the priority of this ACL. If more than one ACL is applied to an interface, then the match criteria for the highest sequence ACLs are checked first. A lower number indicates higher priority. If a sequence number is already in use for this interface and direction, the specified access list replaces the currently attached access list using that sequence number. If you do not specify a sequence number, a sequence number that is one greater than the highest sequence number currently in use for this interface and direction is used. The valid range is 1-4294967295.

6.1.4.1 Assigning an ACL to an Interface

1. Open the **ACL Interface Configuration** page.
2. Select the interface from the **Slot/Port** field to which you want to bind the ACL.
3. Select the type of ACL in the **ACL Type** field.
4. Select the ACL ID or name to bind to the interface.

Note...

Whenever an ACL is assigned on a port, LAG, or VLAN, flows from that ingress interface that do not match the ACL are matched to the default rule, which is Drop unmatched packets.

5. Specify the priority in the **Sequence** field.
6. Click **Submit**.

   The ACL is attached to the specified interface(s).

6.1.4.2 Removing an ACL from an Interface

If an ACL is bound to an interface, the **Remove** button appears on the page when you select the interface from the **Slot/Port** dropdown menu. To remove the ACL from the interface, select the type of ACL to remove and its ID or name, and then click **Remove**.

6.1.5 VLAN ACL Configuration

Use this page to configure ACLs to apply to VLANs on your system rather than to ports. At the bottom of the page, the table displays any currently-configured ACLs for the selected VLAN. You can also bind an ACL to a port; see 6.1.4ACL Interface Configuration393 for information.

To display this page, click **QoS > Access Control Lists > VLAN ACL Configuration** in the navigation tree.

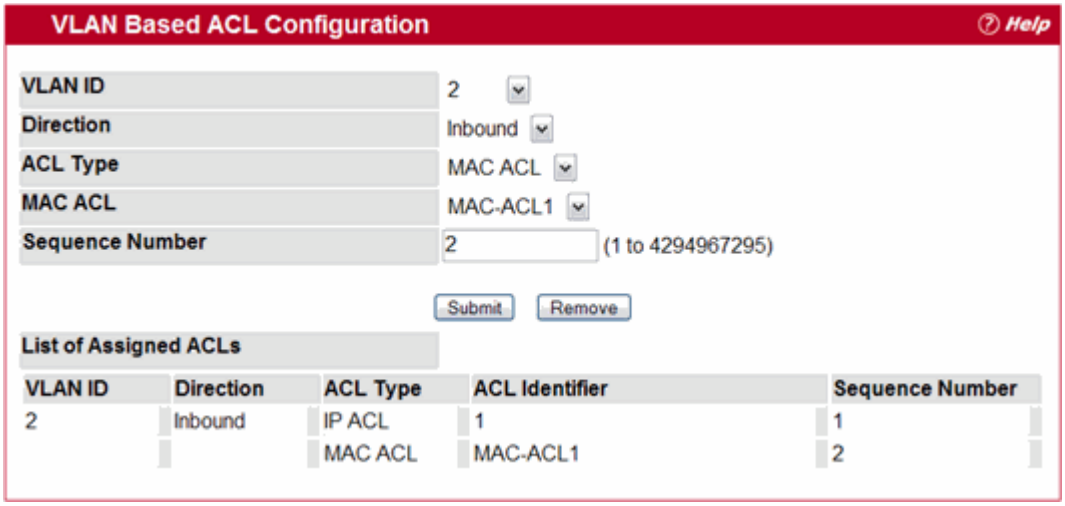

Figure 6-16: VLAN-Based ACL Configuration

The table at the bottom of the page displays any currently configured ACLs on the selected VLAN interface.

Table 6-11: VLAN-Based ACL Configuration

Field	Description
VLAN ID	Select the VLAN ID that you want to associate an ACL to.
Direction	Specifies the packet filtering direction for the ACL. The system supports Inbound and Outbound filtering. Inbound filtering means the system applies the ACL rules to packets as they enter the interface.
ACL Type	Use the menu to select the ACL type to which packets are matched. Packets can be matched to IPv4-, IPv6-, and MAC-based ACLs.
IPv4/IPv6/MAC ACL	Select the ACL of the specified type to apply to the interface from the dropdown menu.
ACL Identifier	Displays the ACL Number or Name identifying the ACL assigned to the selected VLAN and direction.
Sequence Number	Assigns the priority of this ACL. If more than one ACL is applied to an interface, then the match criteria for the highest sequence ACLs are checked first. A lower number indicates higher priority. If a sequence number is already in use for this interface and direction, the specified access list replaces the currently attached access list using that sequence number. If you do not specify a sequence number, a sequence number that is one greater than the highest sequence number currently in use for this interface and direction is used. The valid range is 1-4294967295.

6.1.6 ACL Interface/VLAN Summary

Use this page to view all ports and VLANs to which an ACL has been applied. To access the page, click **QoS > Access Control Lists > ACL Interface/VLAN Summary**.

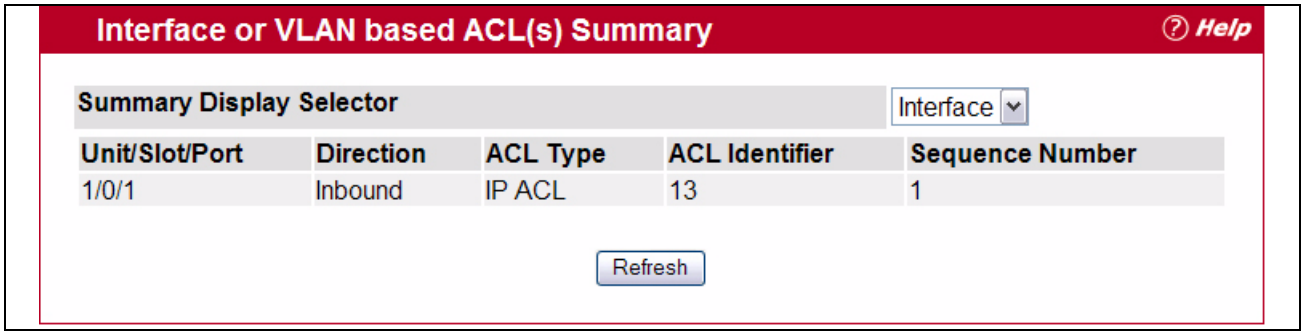

Figure 6-17: Interface/VLAN-Based ACL Configuration

The table at the bottom of the page displays any currently configured ACLs on the selected VLAN interface.

Table 6-12: VLAN-Based ACL Configuration

Field	Description
Summary Display Selector	Select interface or VLAN to display summary. By default summary of Interface-based ACL(s) is displayed.
Slot/Port	Displays the interfaces to which the IP ACL applies.
VLAN ID	Displays the VLAN(s) to which the IP ACL applies.
Direction	The direction of packet traffic affected by the IP ACL. The system supports inbound and outbound filtering.
ACL Type	Displays the type of ACL assigned to selected VLAN and direction.
ACL Identifier	Displays the ACL Number (for IPv4 ACLs) or the ACL Name (for IPv6 and MAC ACLs), which identifies the ACL assigned to the selected VLAN and direction.
Sequence Number	Displays the sequence number signifying the order of specified ACL relative to other ACLs assigned to selected VLAN and direction.

Select **Interface** of **Vlan Id** to display either interface-based or VLAN ID-based ACLs.

6.2 Configuring Differentiated Services

The QoS feature contains Differentiated Services (DiffServ) support that allows traffic to be classified into streams and given certain QoS treatment in accordance with defined per-hop behaviors.

Standard IP-based networks are designed to provide “best effort” data delivery service. “Best effort” service implies that the network delivers the data in a timely fashion, although there is no guarantee that it will. During times of congestion, packets may be delayed, sent sporadically, or dropped. For typical Internet applications, such as e-mail and file transfer, a slight degradation in service is acceptable and in many cases unnoticeable. Conversely, any degradation of service has undesirable effects on applications with strict timing requirements, such as voice or multimedia.

6.2.1 Defining DiffServ

To use DiffServ for QoS, the web pages accessible from the Differentiated Services menu page must first be used to define the following categories and their criteria:

1. Class: Create classes and define class criteria.
2. Policy: Create policies, associate classes with policies, and define policy statements.
3. Service: Add a policy to an inbound interface

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.

Packet processing begins by testing the class match criteria for a packet. A policy is applied to a packet when a class match within that policy is found.

The Differentiated Services menu page contains links to the various Diffserv configuration and display features.

To display the page, click **Quality of Service > Differentiated Services** in the navigation menu. The Differentiated Services menu page contains links to the following features:

- Diffserv Configuration
- Class Configuration
- Class Summary
- Policy Configuration
- Policy Summary
- Policy Class Definition
- Policy Attribute Summary
- Service Configuration
- Service Summary
- Service Statistics
- Service Detailed Statistics

6.2.2 Diffserv Configuration

Packets are filtered and processed based on defined criteria. The filtering 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.

The configuration process begins with defining one or more match criteria for a class. Then one or more classes are added to a policy. Policies are then added to interfaces.

Packet processing begins by testing the match criteria for a packet. The 'all' class type option defines that each match criteria within a class must evaluate to true for a packet to match that class. The 'any' class type option defines that at least one match criteria must evaluate to true for a packet to match that class. Classes are tested in the order in which they were added to the policy. A policy is applied to a packet when a class match within that policy is found.

Use the Diffserv Configuration page to display DiffServ General Status Group information, which includes the current administrative mode setting as well as the current and maximum number of rows in each of the main DiffServ private MIB tables.

To display the page, click **Quality of Service > Differentiated Services > Diffserv Configuration** in the navigation menu.

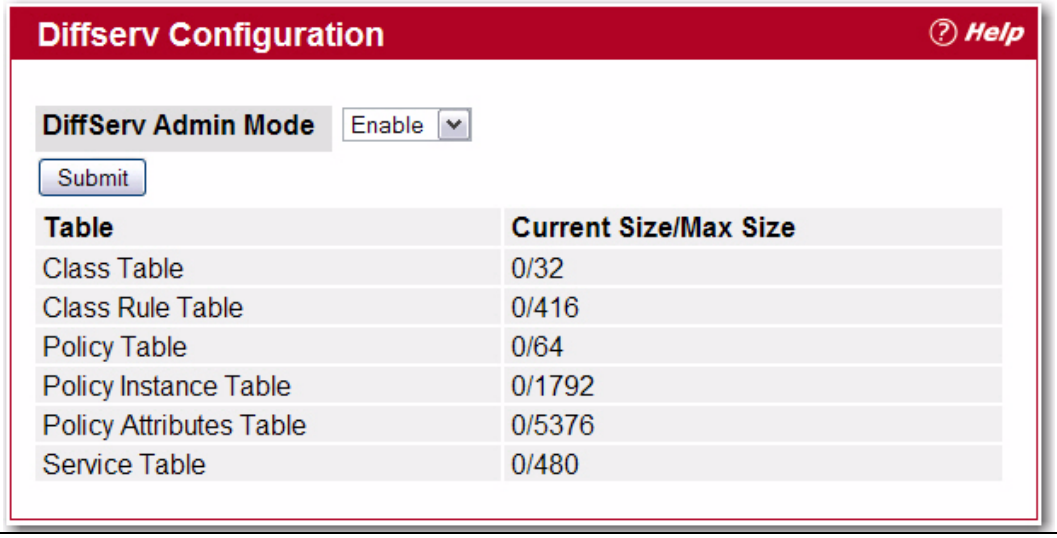

Figure 6-18: Diffserv Configuration

Table 6-13: Diffserv Configuration Fields

Field	Description
Diffserv Admin Mode	Turns admin mode on and off. The default value is **Enable**. While disabled, the DiffServ configuration is retained and can be changed, but it is not active. While enabled, Differentiated Services are active.
MIB Table	
Class Table	Displays the current and maximum number of rows of the class table.
Class Rule Table	Displays the current and maximum number of rows of the class rule table.
Policy Table	Displays the current and maximum number of rows of the policy table.
Policy Instance Table	Displays the current and maximum number of rows of the policy instance table.
Policy Attributes Table	Displays the current and maximum number of rows of the policy attributes table.
Service Table	Displays the current and maximum number of rows of the service table.

If you change the DiffServ admin mode, click **Submit** to apply the change to the system.

6.2.3 Class Configuration

Use the Class Configuration page to add a new Diffserv class name, or to rename or delete an existing class. The page also allows you to define the criteria to associate with a DiffServ class. As packets are received, these DiffServ classes are used to prioritize packets. You can have multiple match criteria in a class. The logic is a Boolean logical AND for this criteria.

To display the page, click **Quality of Service > Differentiated Services > Class Configuration** in the navigation menu.

The fields available on the Class Configuration page depend on whether you create a new class or configure a class that has already been created.

Figure 6-19 shows the Class Configuration page when the Class Selector option is Create.

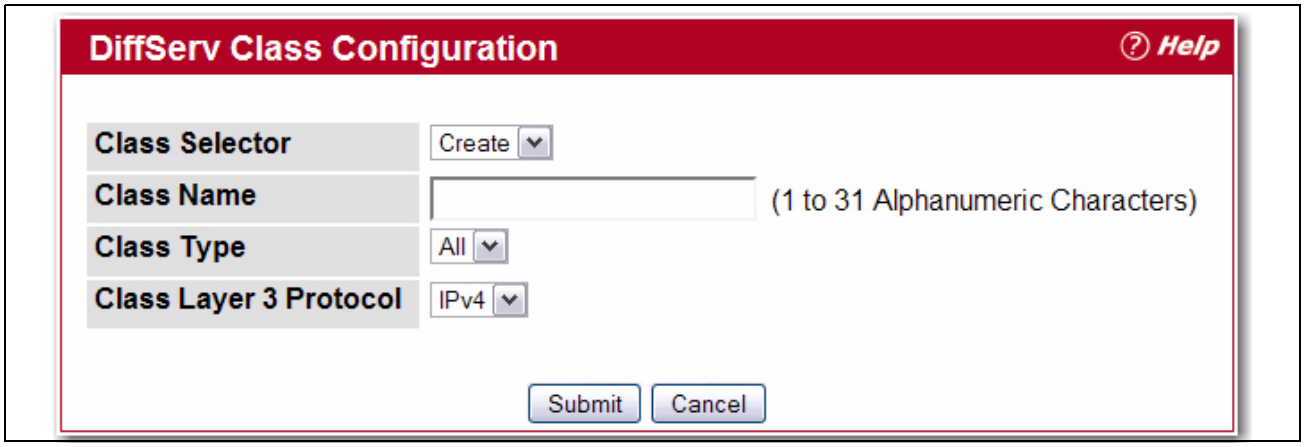

Figure 6-19: Diffserv Class Configuration

Figure 6-20 shows the Class Configuration page when the Class Selector option shows a configured class. The class has three class match selectors configured.

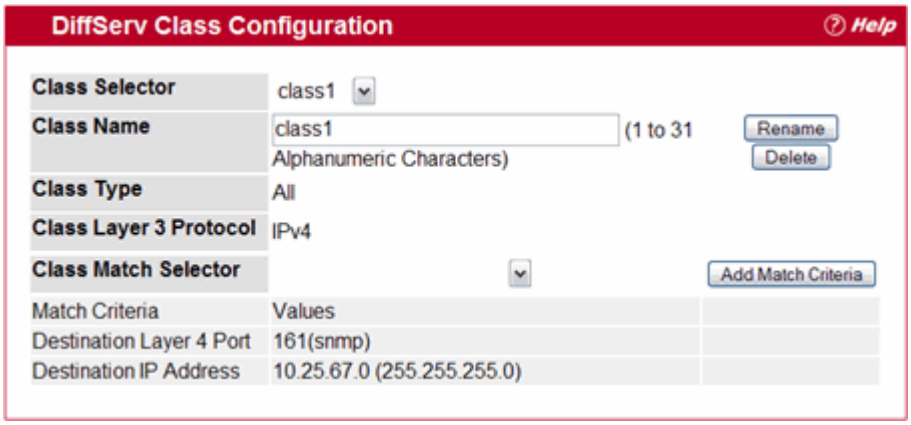

Figure 6-20: Diffserv Class Configuration

Table 6-14: Diffserv Class Configuration Fields

Field	Description
Class Selector	To configure a new DiffServ class, select Create. To modify or view an existing class, select the name of the class from the dropdown menu.
Class Name	Enter a class name. To create a new class, select the class type and click **Submit**. To rename an existing class, click **Rename** after you enter the class name.
Class Type	Lists all of the class types. Currently the hardware supports only the **Class Type** value **All**, which means 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.

Table 6-14: Diffserv Class Configuration Fields (Continued)

Field	Description
Class Layer 3 Protocol	Indicates how to interpret layer 3. This lists types of packets supported by Diffserv. The Layer 3 Protocol option is available only when user selects **All** as **Class Type**. The options are: • IPv4 • IPv6
Class Match Selector (IPv4)	The menu lists all match criteria you can add to a specified class. To configure the criteria, select a match criteria from the list, and then click Add Match Criteria. The screen changes to the criteria configuration page for that class. After you configure the criteria, click **Submit** to apply the criteria to the class and return to the **Class Configuration** Page. To return to the **Class Configuration** page without applying the criteria, click **Cancel**. The match criteria and configurable fields are as follows: • **Class of Service**: Requires a packet's CoS for incoming packets to match the CoS entered here. The valid range is 0-7. • **Secondary Class of Service**: Requires a packet's secondary CoS for incoming packets to match the CoS entered here. The valid range is 0-7. • **Destination IP Address**: Requires a packet's destination port IP address to match the address listed here. In the **IP Address** field, enter a valid destination IP address in dotted decimal format. In the **IP Mask** field, enter a valid subnet mask to determine which bits in the IP address are significant. Note that this is *not* a wildcard mask. • **Destination L4 Port**: Requires a packet's TCP/UDP destination port to match the port you select. Select the desired L4 keyword from the list on which the rule can be based. If you select Other, the screen refreshes and a Port ID field appears. Enter a user-defined Port ID by which packets are matched to the rule. • **Destination MAC Address**: Requires a packet's Destination MAC Address for incoming packets to match the address entered here. In the **MAC Address** field, enter a valid destination MAC address in dotted decimal format. In the **MAC Mask** field, use F's and zeros to configure the MAC mask. An F means that the bit is checked, and a zero in a bit position means that the data is not significant. 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. • **EtherType**: Requires a frames' Ethertype to match the Ethertype listed you select. If you select User Value, you can enter a custom EtherType value. The value you enter specifies a customized Ethertype to compare against an Ethernet frame. The valid range of values is (0x0600 to 0xFFFF). • **Any**: All packets are considered to match the specified class and no additional input information is needed. • **IP DSCP**: Matches the packet's DSCP to the class criteria's when selected. Select the DSCP type from the drop-down menu. or enter a DSCP value to match. If you select Other, enter a custom value in the **DSCP Value** field that appears. • **IP Precedence**: Matches the packet's IP Precedence value to the class criteria's when Enter a value in the range of 0-7. • **IP TOS**: Matches the packet's Type of Service bits in the IP header to the class criteria's when selected and a value is entered. In the **TOS Bits** field, enter a two-digit hexadecimal number to match the bits in a packet's TOS field. In the **TOS Mask** field, specify the bit positions that are used for comparison against the IP TOS field in a packet. • **Protocol**: Requires a packet's layer 4 protocol to match the protocol you select. If you select Other, enter a protocol number in the field that appears. The valid range is 0-255. • **Reference Class**: Selects a class to start referencing for criteria. If the specified class references another class, the Reference Class match criterion disappears from the match list to prevent you adding another class reference, since a specified class can reference at most one other class of the same type. Additionally, a **Remove Class Reference** button appears on the screen. Click the button to remove the current class reference.

Table 6-14: Diffserv Class Configuration Fields (Continued)

Field	Description
Class Match Selector (cont.) (IPv4)	• **Source IP Address**: Requires a packet's source port IP address to match the address listed here. In the **IP Address** field, enter a valid source IP address in dotted decimal format. In the **IP Mask** field, enter a valid subnet mask to determine which bits in the IP address are significant. Note that this is *not* a wildcard mask. • **Source L4 Port**: Requires a packet's TCP/UDP source port to match the port you select. Select the desired L4 keyword from the list on which the rule can be based. If you select Other, the screen refreshes and a Port ID field appears. Enter a user-defined Port ID by which packets are matched to the rule. • **Source MAC Address**: Requires a packet's source MAC Address for incoming packets to match the address entered here. In the **MAC Address** field, enter a valid source MAC address in dotted decimal format. In the **MAC Mask** field, use F's and zeros to configure the MAC mask. An F means that the bit is checked, and a zero in a bit position means that the data is not significant. 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. • **VLAN**: Requires a packet's VLAN ID for incoming packets to match the VLAN ID you enter. • **Secondary VLAN**: Requires a packet's VLAN ID for incoming packets to match the VLAN ID you enter.
Class match Selector (IPv6-only. For other fields not listed here, see the description in the IPv4 list above.	• **Destination IPv6 Address**: Requires a packet's destination IPv6 address to match the address listed here. Enter a valid destination IP address in dotted decimal format. • **Destination L4 Port**: Requires a packet's TCP/UDP destination port to match the port you select. Select the desired L4 keyword from the list on which the rule can be based. If you select Other, the screen refreshes and a Port ID field appears. Enter a user-defined Port ID by which packets are matched to the rule. • **Any**: All packets are considered to match the specified class and no additional input information is needed. • **Flow Label**: Flow label is 20-bit number that is unique to an IPv6 packet, used by end stations to signify quality-of-service handling in routers. The flow label of the incoming packet must match this value. • **IP DSCP**: Matches the packet's DSCP to the class criteria's when selected. Select the DSCP type from the drop-down menu. or enter a DSCP value to match. If you select Other, enter a custom value in the **DSCP Value** field that appears. • **Protocol**: Requires a packet's layer 4 protocol to match the protocol you select. If you select Other, enter a protocol number in the field that appears. The valid range is 0-255. • **Reference Class**: Selects a class to start referencing for criteria. If the specified class references another class, the Reference Class match criterion disappears from the match list to prevent you adding another class reference, since a specified class can reference at most one other class of the same type. Additionally, a **Remove Class Reference** button appears on the screen. Click the button to remove the current class reference. • **Source IPv6 Address**: Requires a packet's source port IPv6 address to match the address listed here. Enter a valid source IP address in dotted decimal format. • **Source L4 Port**: Requires a packet's TCP/UDP source port to match the port you select. Select the desired L4 keyword from the list on which the rule can be based. If you select Other, the screen refreshes and a Port ID field appears. Enter a user-defined Port ID by which packets are matched to the rule.

6.2.4 Class Summary

Use the Class Summary page to view the DiffServ classes configured on the system. The table on the page also indicates whether a class references another class.

To display the page, click **Quality of Service > Differentiated Services > Class Summary** in the navigation menu.

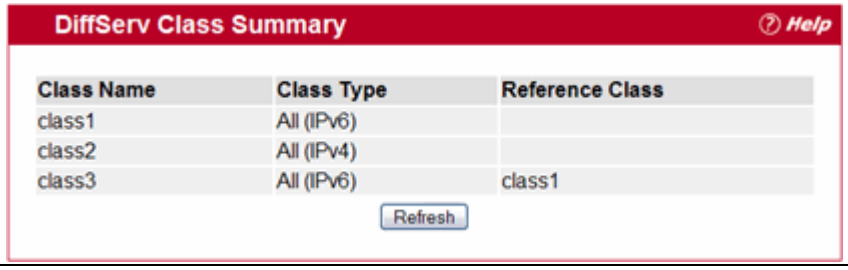

Figure 6-21: Class Summary

Table 6-15: Class Summary Fields

Field	Description
Class Name	Shows the user-defined name of the DiffServ class.
Class Type	Currently the hardware supports only the class type value **All**, which means 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.
Reference Class	If the class is configured with a Reference Class as part of the match criteria, this column shows the name of the configured class that is referenced.

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

6.2.5 Policy Configuration

Use the Policy Configuration page to associate a collection of classes with one or more policy statements.

To display the page, click **Quality of Service > Differentiated Services > Policy Configuration** in the navigation menu.

The fields available on the Policy Configuration page depend on whether you create a new class or configure a class that has already been created.

Figure 6-22 shows the Policy Configuration page when the Policy Selector option is Create.

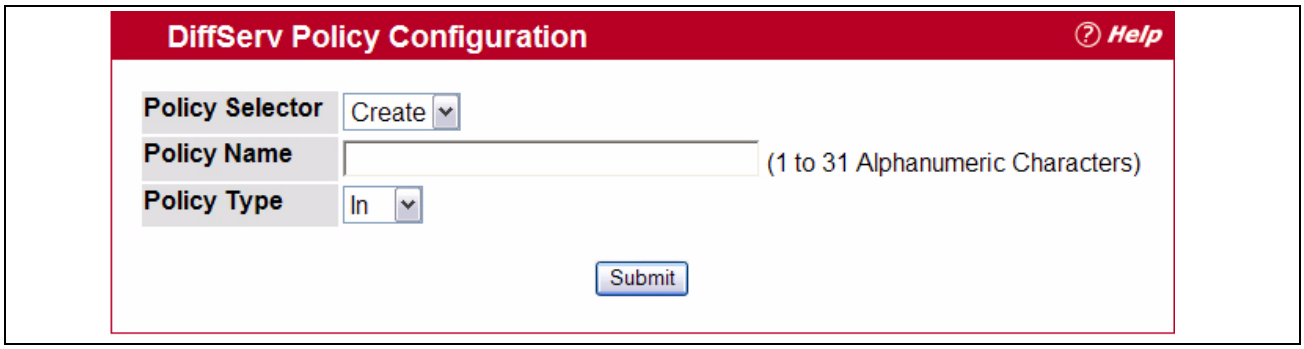

Figure 6-22: Policy Configuration

Figure 6-23 shows the Policy Configuration page when the Policy Selector option shows a configured policy that has a member class. To configure a member class, see 6.2.3Class Configuration399.

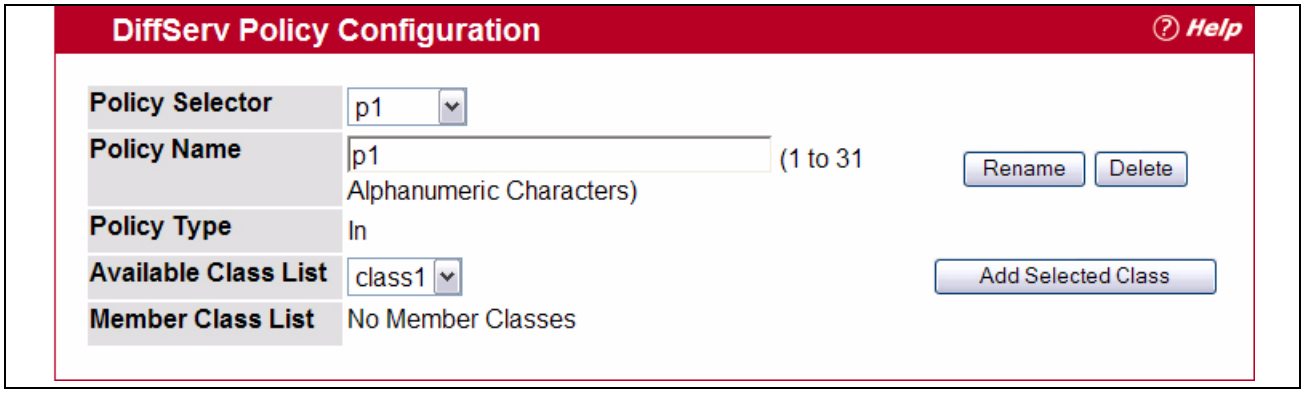

Figure 6-23: Policy Configuration

Table 6-16: Policy Configuration Fields

Field	Description
Policy Selector	To create a new policy, select Create from the menu; another page appears to facilitate creation of a new policy. To change a policy name or to modify the class list members, select the policy name from the menu. To delete an existing policy select it from the menu, and then click **Delete**.

Table 6-16: Policy Configuration Fields (Continued)

Field	Description
Policy Name	If you select Create from the Policy Selector menu, enter a name to associate with the class(es). The name is a case-sensitive alphanumeric string from 1 to 31 characters uniquely identifying a policy. To modify the name of an existing policy, select it from the Policy Selector menu and enter a new name in the **Policy Name** field, and then click **Rename**.
Policy Type	The available policy type is *In*, which indicates the type is specific to inbound traffic. *Out* indicates the type is specific to outbound traffic direction. This field is only configurable when you create a new policy. After policy creation, this becomes a non-configurable field displaying the configured policy type.
Available Class List	The menu lists all existing DiffServ class names. The list is automatically updated as a new class is added or removed from the policy. To associate a DiffServ class with a policy, select the name of the class from the list, and then click **Add Selected Class**.
Member Class List	The menu lists all DiffServ classes that have been added to the policy. names. To remove a DiffServ class from a policy, select the name of the class from the list, and then click **Remove Selected Class**. This list is automatically updated as a new class is added or removed from the policy.

6.2.6 Policy Summary

Use the Policy Summary page to view the DiffServ policies and their associated classes that are configured on the system.

To display the page, click **Quality of Service > Differentiated Services > Policy Summary** in the navigation menu.

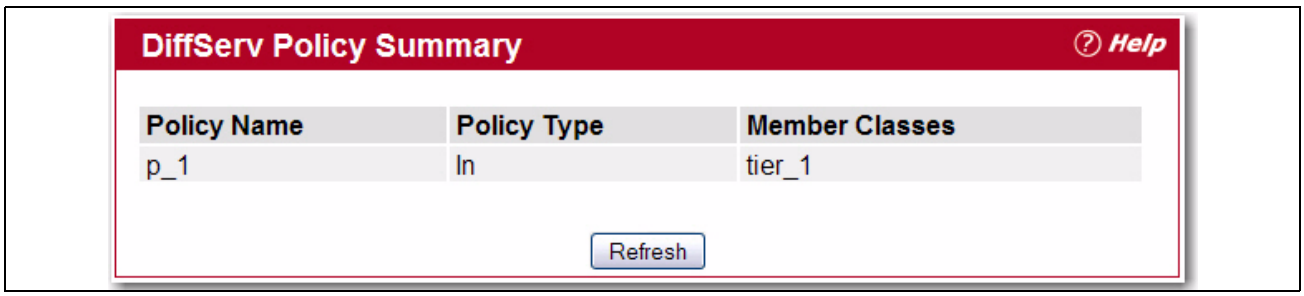

Figure 6-24: Policy Summary

Table 6-17: Policy Summary Fields

Field	Description
Policy Name	Shows the user-defined name of the DiffServ policy.
Policy Type	Displays the policy type value **In**, which means the type is specific to inbound traffic, or **Out**, which means the type is specific to Outbound traffic.
Member Class	Displays name of each class instance within the policy.

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

6.2.7 Policy Class Definition

Use the Policy Class Definition page to associate a class to a policy and to define attributes for that policy-class instance.

To display the page, click **Quality of Service > Differentiated Services > Policy Class Definition** in the navigation menu.

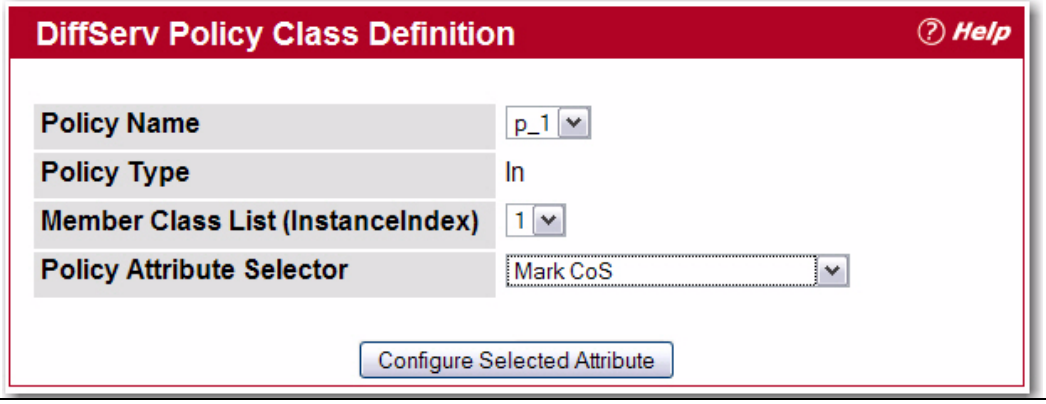

Figure 6-25: Policy Class Definition

Table 6-18 describes all fields available on these pages.

Table 6-18: Policy Class Definition Fields

Field	Description
Policy Name	Select the policy to associate with a member class from the menu.
Policy Type	The read-only field shows the type of policy.
Member Class List	Select the member class to associate with this policy name from the menu.

Policy Attribute Selector	The menu lists all attributes supported for this type of policy, from which one can be selected. To configure the attributes, select an attribute from the list, and then click **Configure Selected Attribute**. The screen changes to the attribute configuration page for that attribute. After you configure the attribute, click **Submit** to apply the criteria to the class and return to the **Policy Class Definition** page. To return to the **Policy Class Definition** page without applying the attribute, click **Cancel**. The attributes and configurable fields are as follows:
	• **Assign Queue**: Assigns the packets of this policy-class to a queue. Enter an integer from 0-7 in the **Queue Id Value** field.
	• **Drop Packets**: Select this field to drop packets for this policy-class. There are no fields to configure. Once you select Drop, click **Configure Select Attribute**, and then click **Submit**, the attribute is added to the policy.
	• **Mark CoS**: Enter the specified Class of Service queue number to mark all packets for the associated traffic stream with the specified class of service value 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. The CoS value is an integer from 0 to 7.
	• **Mark CoS As Secondary CoS**: This option marks all packets for the associated traffic stream with the specified secondary class of service value (the inner 802.1Q tag of a double VLAN tagged packet) in the priority field of the 802.1p header. If the packet does not already contain this header, one is inserted. The CoS value is an integer from 0 to 7.
	• **Mark IP DSCP**: Use this attribute to mark all packets for the associated traffic stream with IP DSCP value you choose from the menu.
	• **Mark IP Precedence**: Use this attribute to mark all packets for the associated traffic stream with the IP Precedence value you enter in the **IP Precedence Value** field.
	• **Mirror**: Use this attribute 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. Redirecting and mirroring actions are mutually exclusive
	• **Police Simple**: Use this attribute to establish the traffic policing style for the specified class. The simple form of the police command uses a single data rate and burst size, resulting in two outcomes (conform and violate). The conforming data rate is specified in kilobits-per-second (Kbps) and is an integer from 1 to 4294967295. The conforming burst size is specified in kilobytes (KB) and is an integer from 1 to 128.
	  • **Police Two Rate:** Use this attribute to establish the traffic policing style for the specified class. The two rate form of the policy command uses two data rate and burst size parameters, resulting in three outcomes: conform, exceed and violate. The conforming and peak data rates are specified in kilobits-per-second (kbps) and are integers in the range from 1 to 4294967295. The conforming and peak burst size parameters are specified in kilobytes (KB) and are integers from 1 to 128.
	  • **Redirect Interface:** Use this attribute to apply Redirect Interface to this policy-class and specify the interface or LAG to be used.

Upon selecting a **Policy Attribute** and clicking **Configure Selected Attribute**, a page displays allowing you to configure the selected attribute. Table 6-19 describes the fields available on these pages.

Table 6-19: Attribute Configuration Fields

Field	Description
Policy Name	Displays name of the specified DiffServ policy.
Policy Type	Displays type of the specified policy (**In** or **Out**).
Class Name	Displays name of the DiffServ class to which this policy is attached.
Mark CoS	Used to mark a packet to notify the network of the precedence of the packet. The range is 0 to 7.
IP Precedence Value	Used to mark a packet to notify the network of the importance of (service level associated with) the packet. The range is 0 to 7.
Policing Style	Displays the type of policing based upon the **Policy Attribute** selected.
Slot/Port	Used to specify the redirect interface or LAG when the **Redirect** attribute is selected.
Color Mode	Displays the color mode. The default is Color Blind.
Color Conform Class	Color Aware mode requires the existence of one or more color classes that are valid for use with this policy instance. A valid color class contains a single, non-excluded match criterion for one of the following fields (provided the field does not conflict with the classifier of the policy instance itself): • CoS • IP DSCP • IP Precedence • Secondary CoS
Committed Rate (Kbps)	Used to monitor arrival rate of incoming packets for this class. The range is 1 to 4294967295 kilobits per second (Kbps).
Committed Burst Size (KB)	Used to determine the amount of conforming traffic allowed. The range is 1 to 128 KBytes.
Peak Rate (Kbps)	Used to monitor arrival rate of incoming packets for this class. The range is 1 to 4294967295 kilobits per second (Kbps).

Table 6-19: Attribute Configuration Fields (Continued)

Field	Description
Peak Burst Size (KB)	Used in determining the amount of exceeding traffic allowed. This value must be equal to or greater than committed burst size. The range is 1 to 128 KBytes.
Conform Action	Used to determine the actions taken on packets that conform to the policing metric. The options include: • **Drop** - These packets are immediately dropped. • **Mark CoS** - These packets are marked by DiffServ with the specified CoS value before being presented to the system forwarding element. This selection requires that the Mark CoS value field be set. • **Mark CoS as Secondary CoS** - For double-tagged packets, the 802.1p tag is marked by DiffServ with the (original) 802.1p value of the inner tag before the packet is presented to the system forwarding element. • **Mark IP DSCP** - These packets are marked by DiffServ with the specified DSCP value before being presented to the system forwarding element. This selection requires the DSCP value field to be set. • **Mark IP Precedence** - These packets are marked by DiffServ with the specified IP Precedence value before being presented to the system forwarding element. This selection requires that the **Mark IP Precedence** value field be set. • **Send** - These packets are presented unmodified by DiffServ to the system forwarding element. The default is **Send**.
Exceed Action	The action to be taken on excess packets per the policing metrics.
Violate Action	Determines the actions to be taken on packets that violate the policing metric. The options include: • **Drop** - These packets are immediately dropped. • **CoS** - These packets are marked by DiffServ with the specified CoS value before being presented to the system forwarding element. This selection requires that the Mark CoS value field be set. • **CoS as Secondary CoS** - For double-tagged packets, the 802.1p tag is marked by DiffServ with the (original) 802.1p value of the inner tag before the packet is presented to the system forwarding element. This allows some of the QoS characteristics of the original packet to be conveyed to switches downstream. • **Mark IP DSCP** - These packets are marked by DiffServ with the specified DSCP value before being presented to the system forwarding element. This selection requires the DSCP value field to be set. • **Mark IP Precedence** - These packets are marked by DiffServ with the specified IP Precedence value before being presented to the system forwarding element. This selection requires that the **Mark IP Precedence** value field be set. • **Send** - These packets are presented unmodified by DiffServ to the system forwarding element. The default is **Drop**.

6.2.8 Policy Attribute Summary

Use the Policy Attribute Summary page to view the attributes associated with the DiffServ policies and their classes.

To display the page, click **Quality of Service > Differentiated Services > Policy Attribute Summary** in the navigation menu.

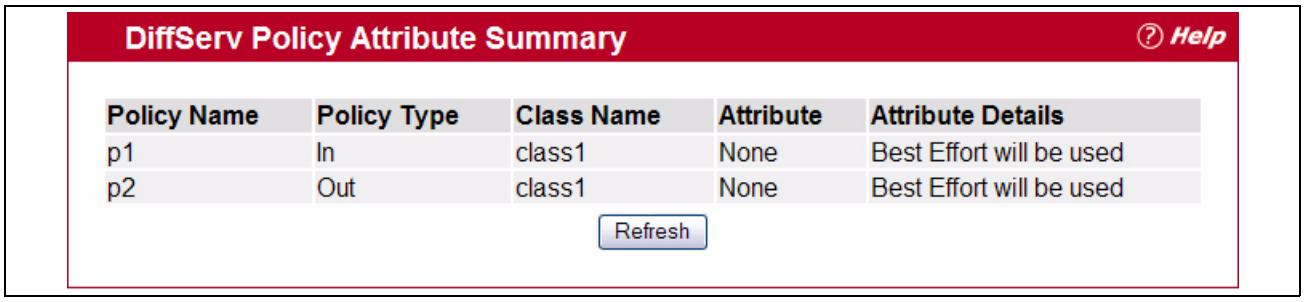

Policy Name	Policy Type	Class Name	Attribute	Attribute Details
p1	In	class1	None	Best Effort will be used
p2	Out	class1	None	Best Effort will be used

Figure 6-26: Policy Attribute Summary

Table 6-20: Policy Attribute Summary Fields

Field	Description
Policy Name	Shows the user-defined name of the DiffServ policy.
Policy Type	Currently the hardware supports only the policy type value **In**, which means the type is specific to inbound traffic, or **Out**, which means the type is specific to Outbound traffic.
Class Name	Displays name of the DiffServ class to which this policy is attached.
Attribute	Displays the attributes attached to the policy class instances.
Attribute Details	Displays the configured values of the attached attributes.

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

6.2.9 Service Configuration

Use the Service Configuration page to activate a policy on a port.

To display the page, click **Quality of Service > Differentiated Services > Service Configuration** in the navigation menu.

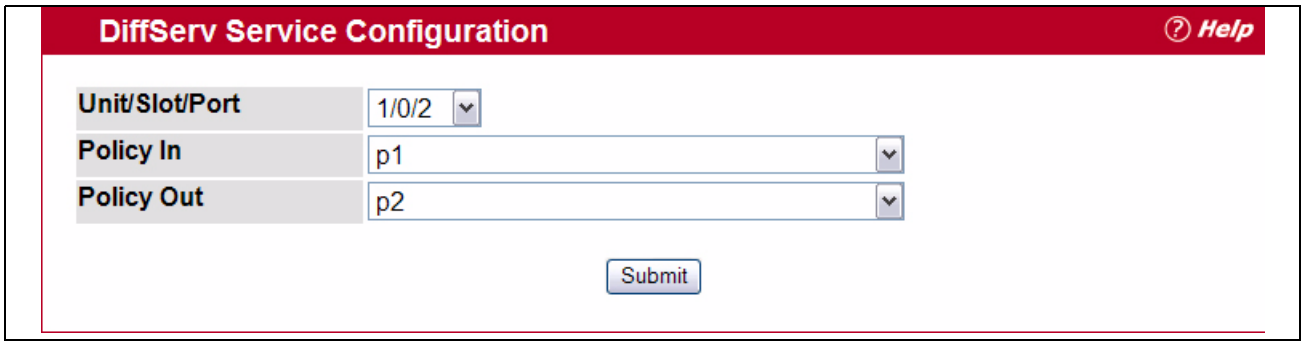

Figure 6-27: Service Configuration

Table 6-21: Service Configuration Fields

Field	Description
Slot/Port	Selects the interface (physical, LAG, or All) to be affected from dropdown menus. This is a list of all valid slot number and port number combinations in the system. For Read/Write users where 'All' appears in the list, select it to specify all interfaces.
Policy In	This lists all the policy names of type 'In' to be associated with the port which can be selected from a dropdown menu. If 'None' is selected, this will detach the policy from the interface in this direction. This field is not shown for Read/Write users where inbound service policy attachment is not supported by the platform.
Policy Out	This lists all the policy names of type 'Out' from which one can be selected. If 'None' is selected, this will detach the policy from the interface in this direction. This field is not shown for Read/Write users where outbound service policy attachment is not supported by the platform.
The following fields display when Slot/Port is specified as 'All'.	
Direction	This selection is only available to Read/Write users when Slot/Port is specified as 'All'. Select the traffic direction of this service interface.
Operational Status	Shows the operational status of this service interface, which is either Up or Down.
Policy Name	Shows the policy associated with the interface.

To activate a policy on an interface, select the interface and the policy, and then click **Submit**.

6.2.10 Service Summary

Use the Service Summary page to display information about the policies activated on a particular interface.

To display the page, click **Quality of Service > Differentiated Services > Service Summary** in the navigation menu.

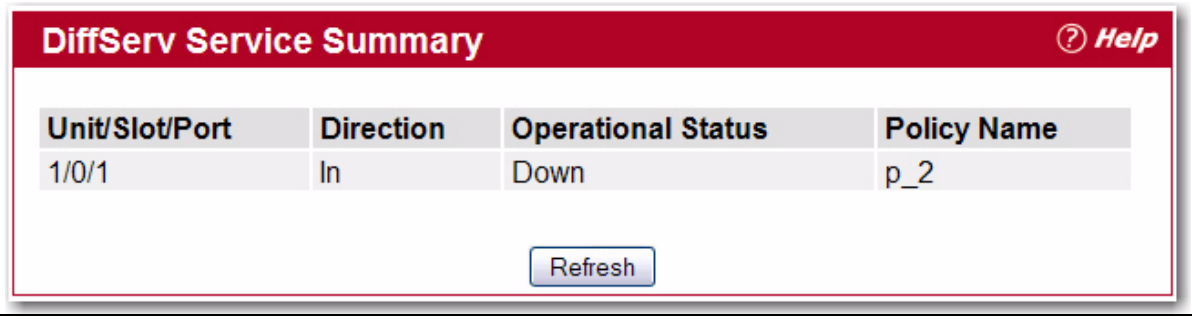

Figure 6-28: Service Summary

Table 6-22: Service Summary Fields

Field	Description
Slot/Port	Shows the interface that has an active policy.
Direction	Shows that the traffic direction of this service interface.
Operational Status	Shows the operational status of this service interface, which is either Up or Down.
Policy Name	Shows the policy associated with the interface.

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

6.2.11 Service Statistics

Use the Service Statistics page to display service-level statistical information about all interfaces that have DiffServ policies attached.

To display the page, click **Quality of Service > Differentiated Services > Service Statistics** in the navigation menu.

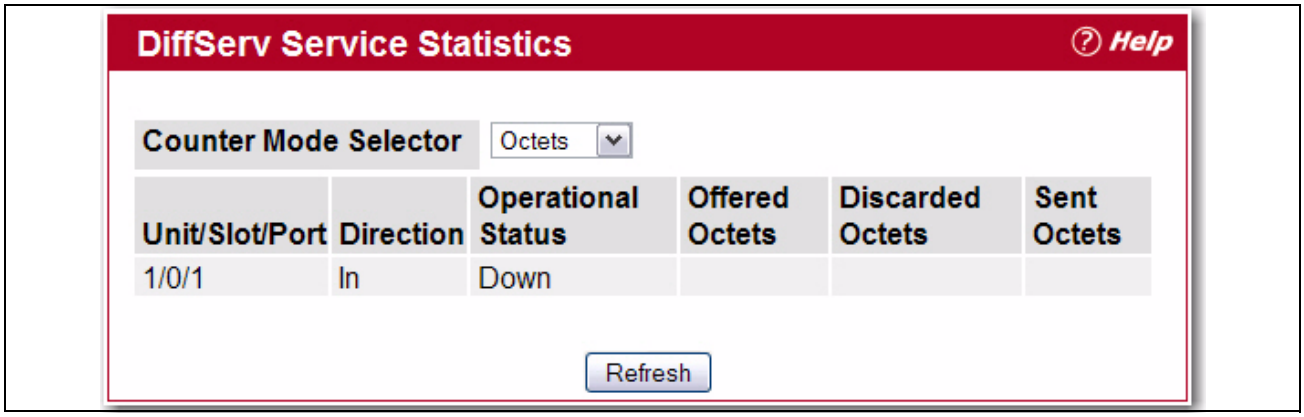

Figure 6-29: Service Statistics

Table 6-23: Service Statistics Fields

Field	Description
Counter Mode Selector	Use the menu to determine the format of the displayed counter values, which must be either Octets or Packets. The default is Octets.
Service Statistics	
Slot/Port	Shows the interface for which service statistics are to display.
Direction	Shows the direction of packets for which service statistics display.
Operational Status	Shows the operational status of this service interface, which is either Up or Down.
Offered Octets	Shows the total number of packets/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.
Discarded Octets	Shows the total number of packets/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.
Sent Octets	Shows the total number of packets/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 information on the screen.

6.2.12 Service Detailed Statistics

Use the Service Detailed Statistics page to display class-oriented statistical information for the policy, which is specified by the interface and direction. The 'Member Classes' dropdown menu is populated on the basis of the specified interface and direction and hence the attached policy (if any). Highlighting a member class name displays the statistical information for the policy-class instance for the specified interface and direction.

To display the page, click **Quality of Service > Differentiated Services > Service Detailed Statistics** in the navigation menu.

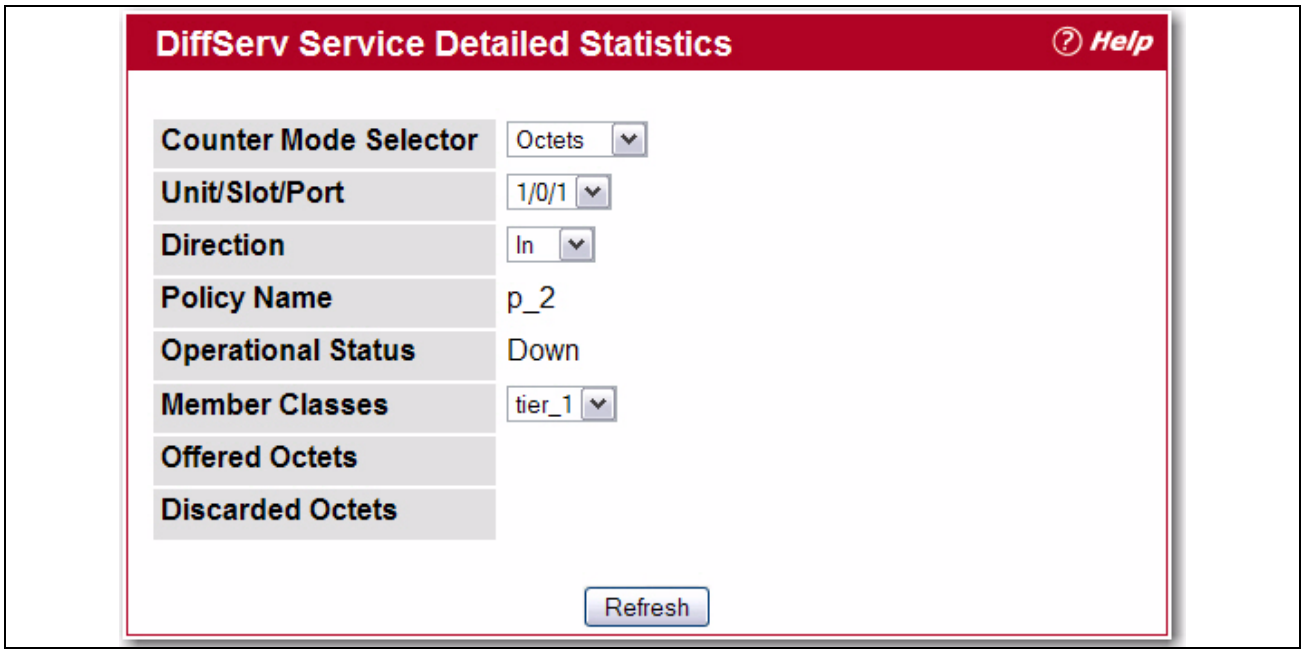

Figure 6-30: Service Detailed Statistics

Table 6-24: Service Detailed Statistics Fields

Field	Description
Counter Mode Selector	Selects the format of the displayed counter values, which must be either Octets or Packets. The default is Octets.
Slot/Port	List of all valid slot number and port number combinations in the system that have a DiffServ policy currently attached in the In direction.
Direction	Selects the direction of packets for which service statistics are to display. Only shows the direction(s) for which a DiffServ policy is currently attached.
Policy Name	Displays the policy currently attached to the selected interface and direction.
Operational Status	Displays the operational status of the policy currently attached to the specified interface and direction. The value is either Up or Down.
Member Classes	List of all DiffServ classes currently defined as members of the selected Policy Name. Choose one member class name at a time to display its statistics. If no class is associated with the chosen policy, then nothing will be populated in the list.
Offered Packets/ Octets	Displays the count of the packets/octets offered to this class instance before the defined DiffServ treatment is applied.
Discarded Packets/Octets	Displays the count of packets/octets discarded for this class instance for any reason due to DiffServ treatment of the traffic class.

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

6.3 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 Diff-Serv 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.

The Class of Service folder contains links to the following features:

- Interface Configuration
- Trust Mode Configuration
- IP DSCP Mapping Configuration
- Interface Queue Configuration
- Interface Queue Status
- Interface Queue Drop Precedence Configuration
- Interface Queue Drop Precedence Status

6.3.1 Interface Configuration

Use the Interface Configuration page to apply an interface shaping rate to all ports or to a specific port.

To display the Interface Configuration page, click **Quality of Service > Class of Service > Interface Configuration** in the navigation menu.

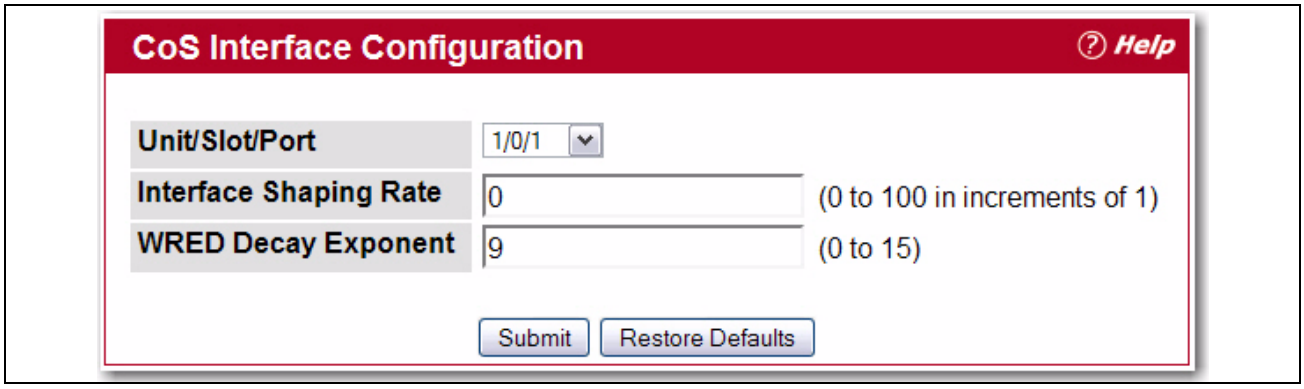

Figure 6-31: Interface Configuration

Table 6-25: Interface Configuration Fields

Field	Description
Slot/Port	Selects the CoS configurable interface to be affected by the Interface Shaping Rate. Select Global to apply a rate to all interfaces. Select an individual port to override the global setting.
Interface Shaping Rate	Sets the limit on how much traffic can leave a port. The limit on maximum transmission bandwidth has the effect of smoothing temporary traffic bursts over time so that the transmitted traffic rate is bounded. The specified value represents a percentage of the maximum negotiated bandwidth. The default value is zero (0). Valid values are 0-100, in increments of 1. A value of 0 means the maximum is unlimited.
WRED Decay Exponent	Specifies the decay exponent value used with the WRED average queue length calculation algorithm. Default value is 9. Valid Range is (0 to 15).

If you make changes to the page, click **Submit** to apply the changes to the system. Click Restore Defaults to reset all interfaces to the default trust value.

6.3.2 Trust Mode Configuration

Use the Trust Mode Configuration page to set the class of service trust mode of an interface. Each port in the switch can be configured to trust one of the packet fields (802.1p or IP DSCP), or to not trust any packet's priority designation (untrusted mode). If the port is set to a trusted mode, it uses a mapping table appropriate for the trusted field being used. This mapping table indicates the CoS queue to which the packet should be forwarded on the appropriate egress port(s). Of course, the trusted field must exist in the packet for the mapping table to be of any use, so there are default actions performed when this is not the case. These actions involve directing the packet to a specific CoS level configured for the ingress port as a whole, based on the existing port default priority as mapped to a traffic class by the current 802.1p mapping table.

Alternatively, when a port is configured as untrusted, it does not trust any incoming packet priority designation and uses the port default priority value instead. All packets arriving at the ingress of an untrusted port are directed to a specific CoS queue on the appropriate egress port(s), in accordance with the configured default priority of the ingress port. This process is also used for cases where a trusted port mapping is unable to be honored, such as when a non-IP packet arrives at a port configured to trust the IP precedence or IP DSCP value.

To display the Trust Mode Configuration page, click **Quality of Service > Class of Service > Trust Mode Configuration** in the navigation menu.

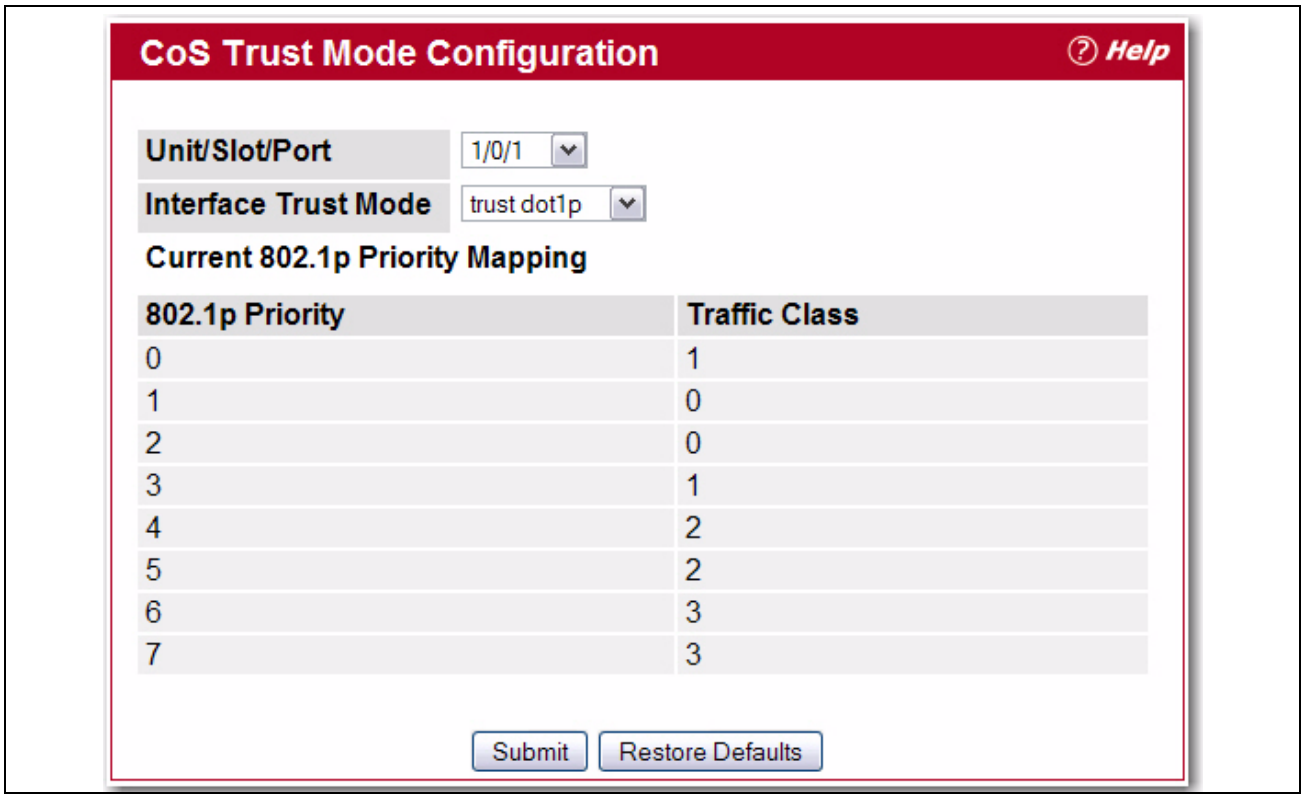

Figure 6-32: Trust Mode Configuration

Table 6-26: Trust Mode Configuration Fields

Field	Description
Slot/Port	The menu contains all CoS configurable interfaces. Select the Global option to apply the same trust mode to all interfaces. Select an individual interface from the menu to override the global settings on a per-interface basis.
Interface Trust Mode	Specifies whether or not an interface (or all interfaces if the Slot/Port field is set to Global) trust a particular packet marking when the packet enters the port. The default value is trust dot1p. The mode can only be one of the following: • untrusted • trust dot1p • trust ip-dscp
Non-IP Traffic Class	This field appears if the trust mode for the selected interface is trust ip-precedence or trust ip-dscp. The field displays the traffic class (queue) to which all non-IP traffic is directed when in the interface trust mode is trust ip-precedence or trust ip-dscp. The valid range is 0 to 6.
Untrusted Traffic Class	This field appears if the trust mode for the selected interface is untrusted. The field displays the traffic class (queue) to which all traffic is directed when in untrusted mode. The valid range is 0 to 6.
Current 802.1p Priority Mapping	Displays the current 802.1p priority mapping configuration.

The **Trust Mode Configuration** page also displays the Current 802.1p Priority Mapping table. For information about 802.1p priority mapping, see Mapping 802.1p Priority269.

- If you make changes to the page, click **Submit** to apply the changes to the system.
- Click **Restore Defaults** to reset the selected interface (or all interfaces, if Global is selected) to the default trust value.

6.3.3 IP DSCP Mapping Configuration

Use the IP DSCP Mapping Configuration page to map an IP DSCP value to an internal traffic class.

To display the IP DSCP Mapping Configuration page, click **Quality of Service > Class of Service > IP DSCP Mapping Configuration** in the navigation menu.

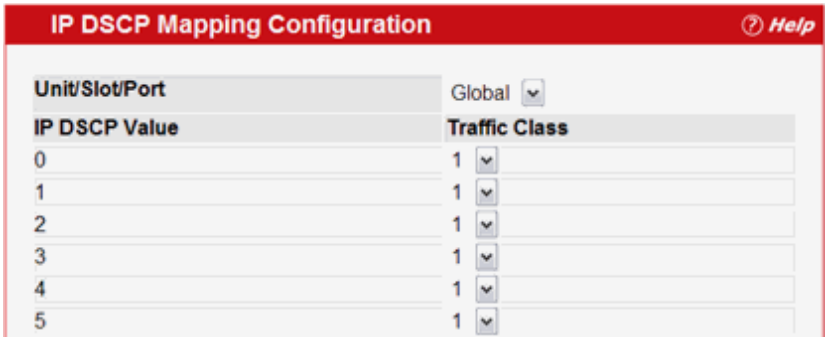

Figure 6-33: IP DSCP Mapping Configuration

Table 6-27: IP DSCP Mapping Configuration Fields

Field	Description
Slot/Port	The menu contains all CoS configurable interfaces. The only option is Global, which means that the IP DSCP mapping configuration applies to all interfaces and cannot be applied on a per-interface basis.
IP DSCP Values	Lists the IP DSCP values to which you can map an internal traffic class. The values range from 0-63.
Traffic Class	The traffic class is the hardware queue for a port. Higher traffic class values indicate a higher queue position. Before traffic in a lower queue is sent, it must wait for traffic in higher queues to be sent. Valid range is 0 to 6.

If you make changes to the page, click **Submit** to apply the changes to the system. Click Restore Defaults to reset all interfaces to the default trust value.

6.3.4 Interface Queue Configuration

Use the Interface Queue Configuration page to define what a particular queue does by configuring switch egress queues. User-configurable parameters control the amount of bandwidth used by the queue, the queue depth during times of congestion, and the scheduling of packet transmission from the set of all queues on a port. Each port has its own CoS queue-related configuration.

The configuration process is simplified by allowing each CoS queue parameter to be configured globally or per-port. A global configuration change is automatically applied to all ports in the system.

To display the Interface Queue Configuration page, click **Quality of Service > Class of Service > Interface Queue Configuration** in the navigation menu.

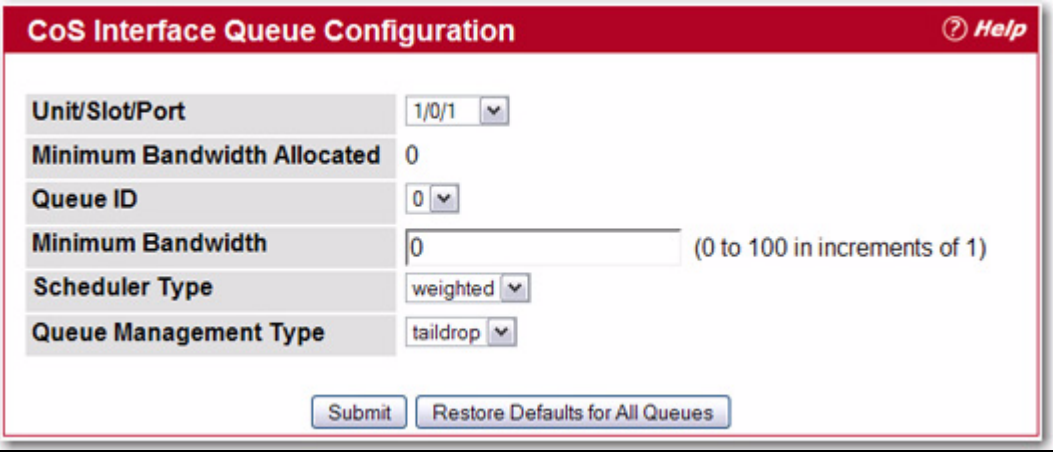

Figure 6-34: Interface Queue Configuration

Table 6-28: Interface Queue Configuration Fields

Field	Description
Slot/Port	Specifies the interface (physical, LAG, or Global) to configure.
Minimum Bandwidth Allocated	Shows the sum of individual Minimum Bandwidth values for all queues in the interface. The sum cannot exceed the defined maximum of 100. This value is considered while configuring the Minimum Bandwidth for a queue in the selected interface.
Queue ID	Use the menu to select the queue per interface to be configured.
Minimum Bandwidth	Specify the minimum guaranteed bandwidth allocated to the selected queue on the interface. Setting this value higher than its corresponding Maximum Bandwidth automatically increases the maximum to the same value. The default value is 0. The valid range is 0 to 100, in increments of 1. The value zero (0) means no guaranteed minimum. The sum of individual Minimum Bandwidth values for all queues in the selected interface cannot exceed defined maximum 100.
Scheduler Type	Selects the type of queue processing from the dropdown menu. Options are **Weighted** and **Strict**. Defining on a per-queue basis allows the user to create the desired service characteristics for different types of traffic. • **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	Displays the type of queue depth management techniques used for all queues on this interface. This is only used if the device supports independent settings per-queue. Queue Management Type can only be Taildrop. The default value is Taildrop. All packets on a queue are safe until congestion occurs. At this point, any additional packets queued are dropped.

- If you make changes to the page, click **Submit** to apply the changes to the system.
- Click **Restore Defaults for all Queues** to reset the settings for the selected interface.
- To reset the defaults for all interfaces, select Global from the **Slot/Port** menu before you click the button.

6.3.5 Interface Queue Status

Use the Interface Queue Status page to view CoS queue interface queue configuration information for each interface.

To display the Interface Queue Status page, click **Quality of Service > Class of Service > Interface Queue Status** in the navigation menu.

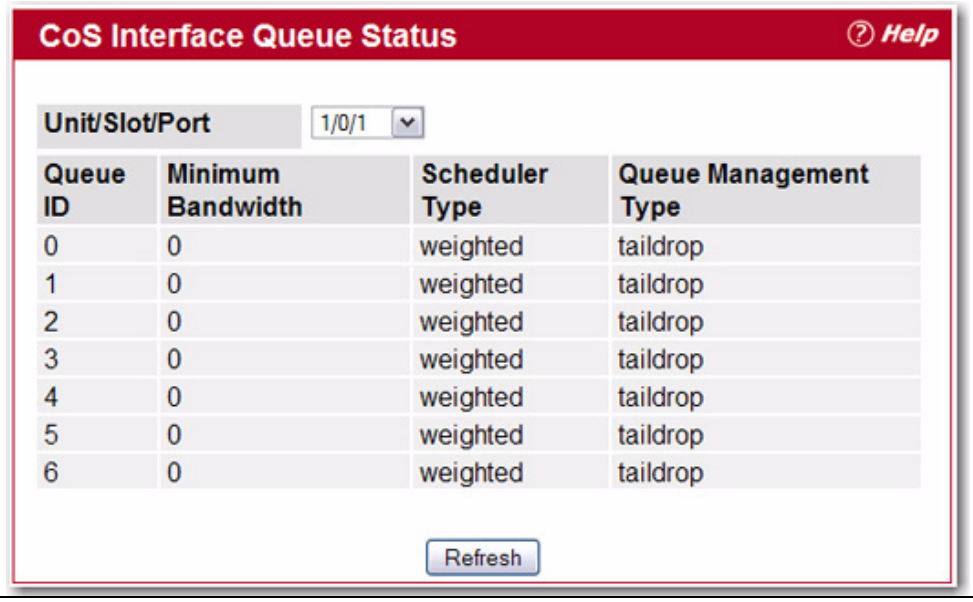

Figure 6-35: Interface Queue Status

Table 6-29: Interface Queue Status Fields

Field	Description
Slot/Port	Select the CoS configurable interface with the queue status information to display.The option Global represents the most recent global configuration settings. These may be overridden on a per-interface basis.
Queue ID	Lists the queues for the interface.
Minimum Bandwidth	Shows the minimum guaranteed bandwidth allotted to this queue on the interface. A value of zero (0) means no guaranteed minimum. The sum of individual Minimum Bandwidth values for all queues in the selected interface cannot exceed the defined maximum of 100.
Scheduler Type	Shows the type of scheduling used for this queue. The available options are as follows: • **Weighted**: Weighted round robin associates a weight to each queue (default). • **Strict**: Strict priority services traffic with the highest priority on a queue first.
Queue Management Type	Displays the queue depth management technique used for queues on this interface, which can only be Taildrop. This is only used if the device supports independent settings per-queue. All packets on a queue are safe until congestion occurs. At this point, any additional packets queued are dropped.

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

6.3.6 Interface Queue Drop Precedence Configuration

Use the **Interface Queue Drop Precedence Configuration** page to configure thresholds for packet loss during times of queue congestion. Each port can have its own drop precedence configuration or all ports can be globally configured.

To display the **Interface Queue Drop Precedence Configuration** page, click **Quality of Service > Class of Service > Interface Queue Drop Precedence Configuration** in the navigation menu.

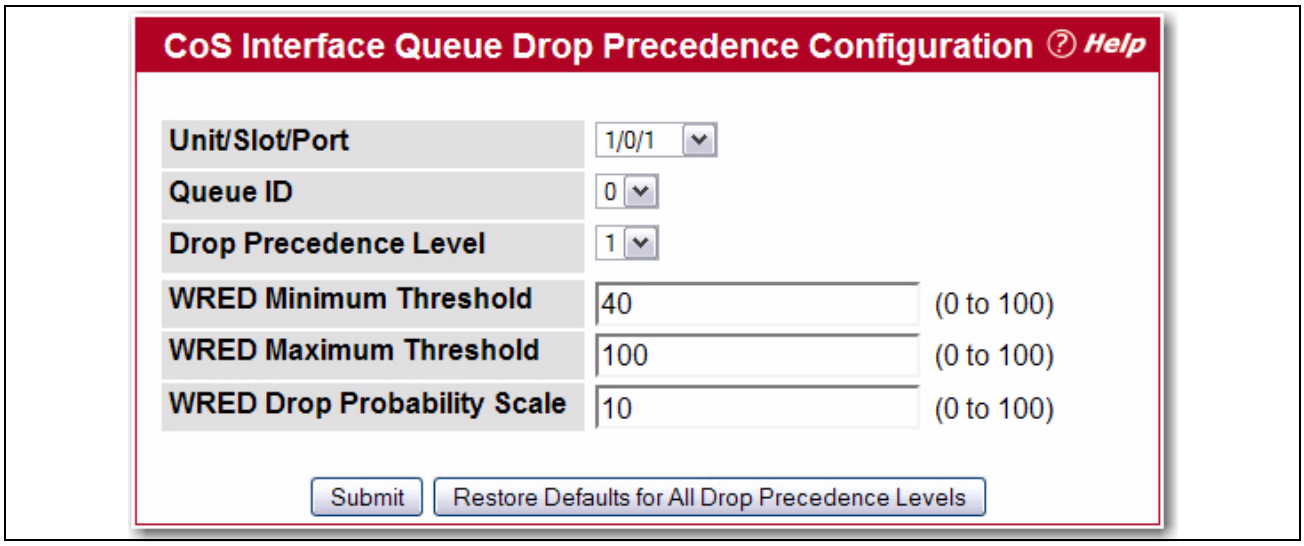

Figure 6-36: Interface Queue Drop Precedence Configuration

Table 6-30: Interface Queue Drop Precedence Configuration Fields

Field	Description
Slot/Port	Specify the interface (physical, LAG, or Global) to configure.
Queue ID	Select a queue to associate with the interface to be configured.

Table 6-30: Interface Queue Drop Precedence Configuration Fields (Continued)

Field	Description
Drop Precedence Level	Select a drop precedence levels (platform-based).
WRED Minimum Threshold	Specify the weighted RED minimum queue threshold below which no packets are dropped for the current drop precedence level. Default values are: • 40/30/20/100 for Drop Precedence Levels 1/2/3/4 on QueueID 0 • 40/30/20/100 for Drop Precedence Levels 1/2/3/4 on QueueID 1 • 40/30/20/100 for Drop Precedence Levels 1/2/3/4 on QueueID 2 • 40/30/20/100 for Drop Precedence Levels 1/2/3/4 on QueueID 3 • 40/30/20/100 for Drop Precedence Levels 1/2/3/4 on QueueID 4 • 40/30/20/100 for Drop Precedence Levels 1/2/3/4 on QueueID 5 • 40/30/20/100 for Drop Precedence Levels 1/2/3/4 on QueueID 6 Valid Range is (0 to 100) in sixteenths of the overall device queue size.
WRED Maximum Threshold	Specify the weighted RED maximum queue threshold above which all packets are dropped for the current drop precedence level. Default values are: • 100/90/80/100 for Drop Precedence Levels 1/2/3/4 on QueueID 0 • 100/90/80/100 for Drop Precedence Levels 1/2/3/4 on QueueID 1 • 100/90/80/100 for Drop Precedence Levels 1/2/3/4 on QueueID 2 • 100/90/80/100 for Drop Precedence Levels 1/2/3/4 on QueueID 3 • 100/90/80/100 for Drop Precedence Levels 1/2/3/4 on QueueID 4 • 100/90/80/100 for Drop Precedence Levels 1/2/3/4 on QueueID 5 • 100/90/80/100 for Drop Precedence Levels 1/2/3/4 on QueueID 6 Valid Range is (0 to 100) in sixteenths of the overall device queue size.
WRED Drop Probability	Specify the packet drop probability for the current drop precedence level. Default value is 10. Valid Range is (0 to 100).

- If you make changes to the page, click **Submit** to apply the changes to the system.
- Click **Restore Defaults for all Queues** to reset the settings for the selected interface.
- To reset the defaults for all interfaces, select Global from the **Slot/Port** menu before you click the button.

6.3.7 Interface Queue Drop Precedence Status

Use the Interface Queue Status page to view interface queue drop precedence configuration information for the CoS queues for each interface.

To display the **Interface Queue Drop Precedence Status** page, click **Quality of Service > Class of Service > Interface Queue Drop Precedence Status** in the navigation menu.

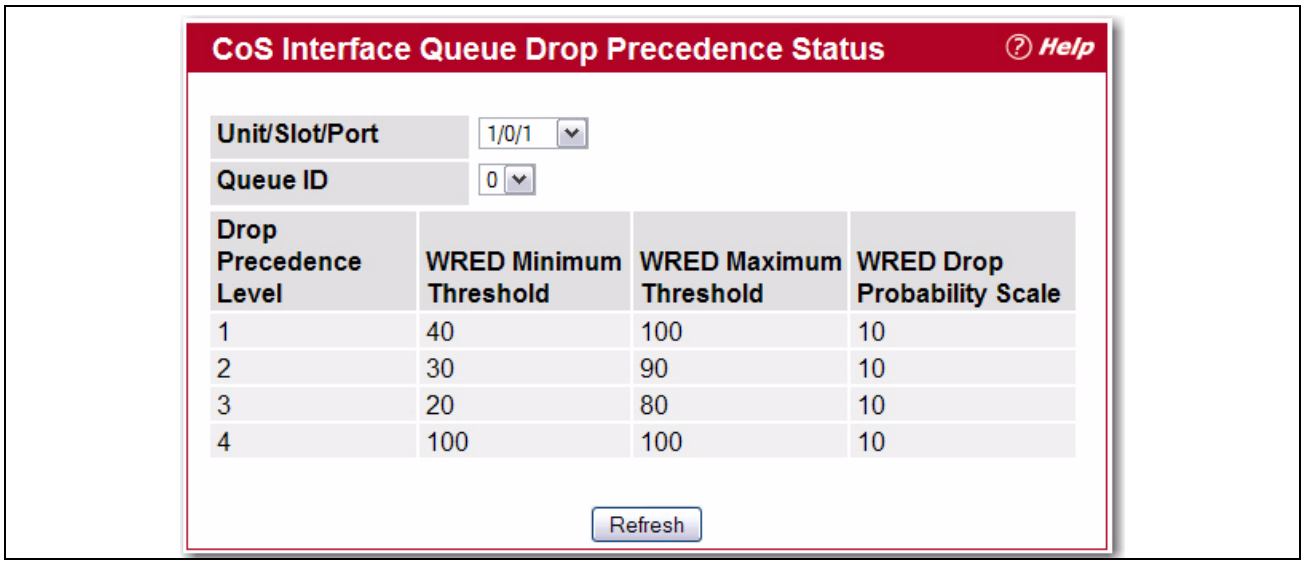

Drop Precedence Level	WRED Minimum Threshold	WRED Maximum Threshold	WRED Drop Probability Scale
1	40	100	10
2	30	90	10
3	20	80	10
4	100	100	10

Figure 6-37: Interface Queue Drop Precedence Status

See 6.3.6Interface Queue Drop Precedence Configuration420 for a description of these fields.

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

6.4 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.

The Auto VoIP folder contains links to the following features:

- Auto VoIP Configuration
- Auto VoIP Summary

6.4.1 Auto VoIP Configuration

Use the Auto VoIP Configuration page to configure the Auto VoIP settings.

To display the Auto VoIP Configuration page, click **Quality of Service > Auto VoIP > Auto VoIP Configuration** in the navigation menu.

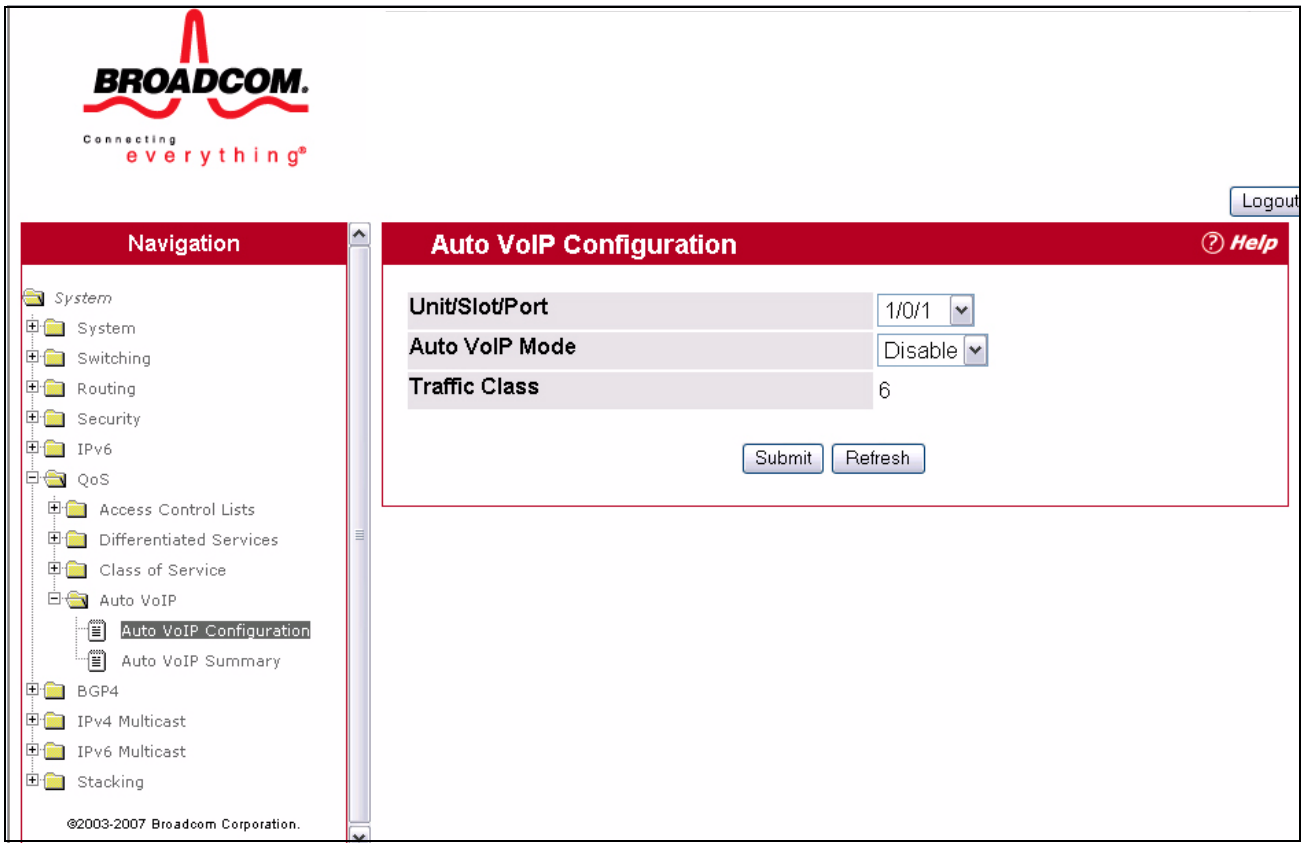

Figure 6-38: Auto VoIP Configuration

Table 6-31: Auto VoIP Configuration Fields

Field	Description
Slot/Port	Specifies all Auto VoIP configurable interfaces. The **All** option represents the most recent configuration settings done for all ports. These settings may be overridden on a per-interface basis.
Auto VoIP Mode	Use to either **Enable** or **Disable** the Auto VoIP mode. The default is **Disable**.
Traffic Class	Displays the traffic class used for VoIP traffic.

- If you change any of the settings on the page, click **Submit** to send the updated configuration to the switch. Configuration changes take effect immediately. These changes will not be retained across a power cycle unless a Save is performed.
- Click **Refresh** to update the page with the most current data from the switch.

6.4.2 Auto VoIP Summary

Use the Auto VoIP Summary page to display the Auto VoIP settings.

To display the Auto VoIP Summary page, click **Quality of Service > Auto VoIP > Auto VoIP Summary** in the navigation menu. A portion of the web page is shown below.

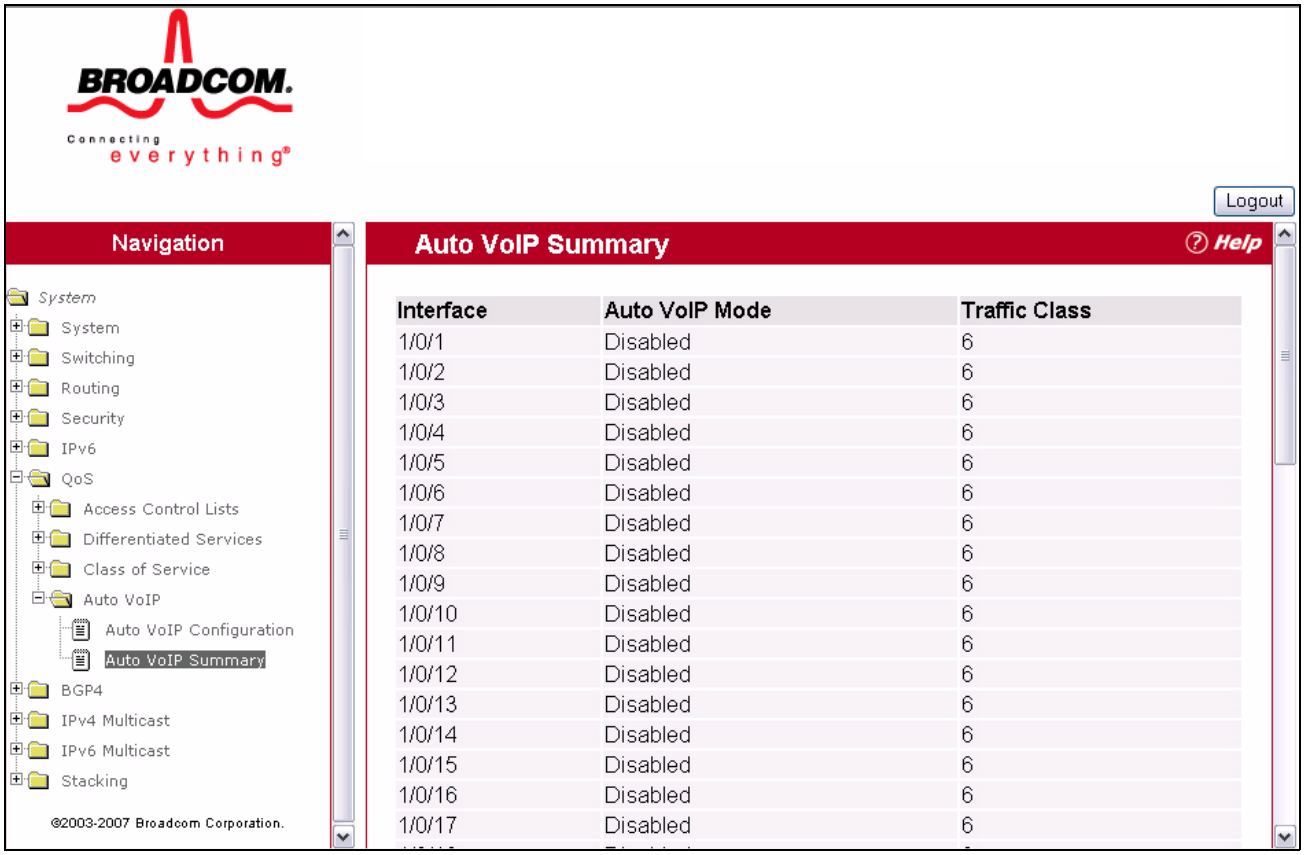

Interface	Auto VoIP Mode	Traffic Class
1/0/1	Disabled	6
1/0/2	Disabled	6
1/0/3	Disabled	6
1/0/4	Disabled	6
1/0/5	Disabled	6
1/0/6	Disabled	6
1/0/7	Disabled	6
1/0/8	Disabled	6
1/0/9	Disabled	6
1/0/10	Disabled	6
1/0/11	Disabled	6
1/0/12	Disabled	6
1/0/13	Disabled	6
1/0/14	Disabled	6
1/0/15	Disabled	6
1/0/16	Disabled	6
1/0/17	Disabled	6

Figure 6-39: Auto VoIP Summary

Table 6-32: Auto VoIP Summary Fields

Field	Description
Interface	Displays the list of Auto VoIP configurable ports available on the switch.
Auto VoIP Mode	Displays whether the mode is enabled or disabled.
Traffic Class	Displays the traffic class used for VoIP traffic.

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

www.kontron.com

6.5 Configuring iSCSI Optimization

iSCSI Optimization provides a means of giving traffic between iSCSI initiator and target systems special Quality of Service (QoS) treatment. This is accomplished by monitoring traffic to detect packets used by iSCSI stations to establish iSCSI sessions and connections. Data from these exchanges is used to create classification rules that assign the traffic between the stations to a configured traffic class. Packets in the flow are queued and scheduled for egress on the destination port based on these rules.

The iSCSI Optimization folder contains links to the following web pages:

- iSCSI Global Configuration
- iSCSI Targets Table
- iSCSI Sessions
- iSCSI Sessions Detailed

6.5.1 iSCSI Global Configuration

Use the iSCSI Optimization-Global Parameters page to configure iSCSI Optimization on the switch.

To display the iSCSI Optimization-Global Parameters page, click **Quality of Service > iSCSI > Global Configuration** in the navigation menu.

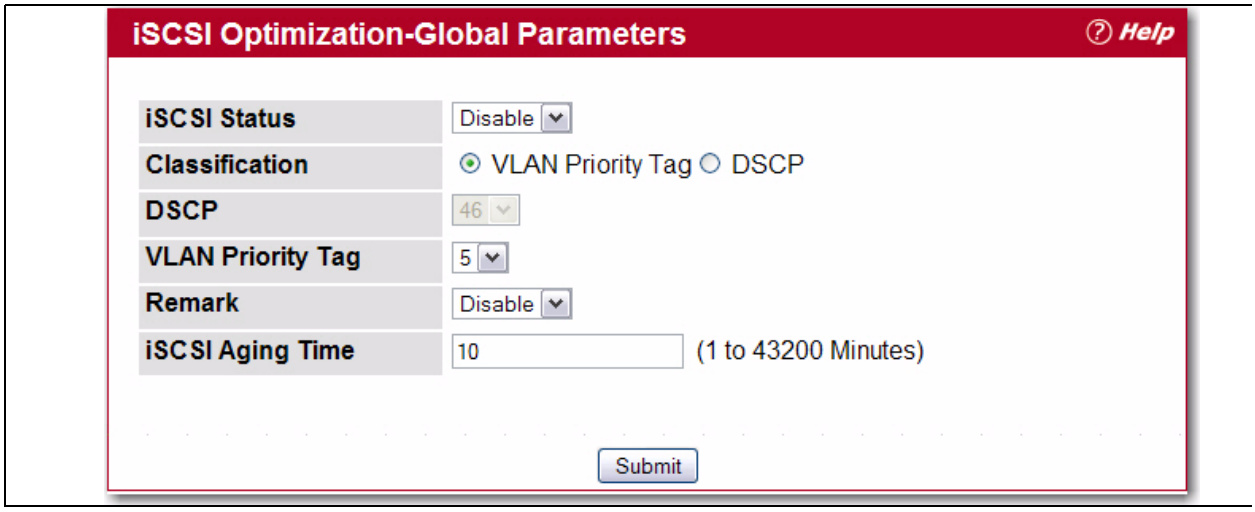

Figure 6-40: iSCSI Optimization-Global Parameters

Table 6-33: iSCSI Optimization Global Parameter Fields

Field	Description
ISCSI Status	Use to either **Enable** or **Disable** iSCSI Optimization. The default is **Disable**.
Classification	Select one of the traffic mapping configuration options: • **VLAN Priority Tag** • **DSCP**

Table 6-33: iSCSI Optimization Global Parameter Fields

Field	Description
DSCP	If using DSCP, assign a DSCP value to iSCSI session packets.
VLAN Priority Tag	If using VLAN Priority, assign a VLAN Priority value to iSCSI session packets.
Remark	Use to either **Enable** or **Disable** the **Remark** mode. The default is **Disable**. Enabling Remark allows the packets to be updated with IEEE 802.1p or IP-DSCP values. Remarking packets with priority data provides special QoS treatment as the packets continue through the network.
iSCSI Aging Time	Set the number of minute a session can be inactive prior to removal.

6.5.2 iSCSI Targets Table

Use the iSCSI Targets Table page to assign target ports/port IP address combinations for iSCSI Optimization on the switch.

To display the iSCSI Optimization-Global Parameters page, click **Quality of Service > iSCSI > Targets** in the navigation menu.

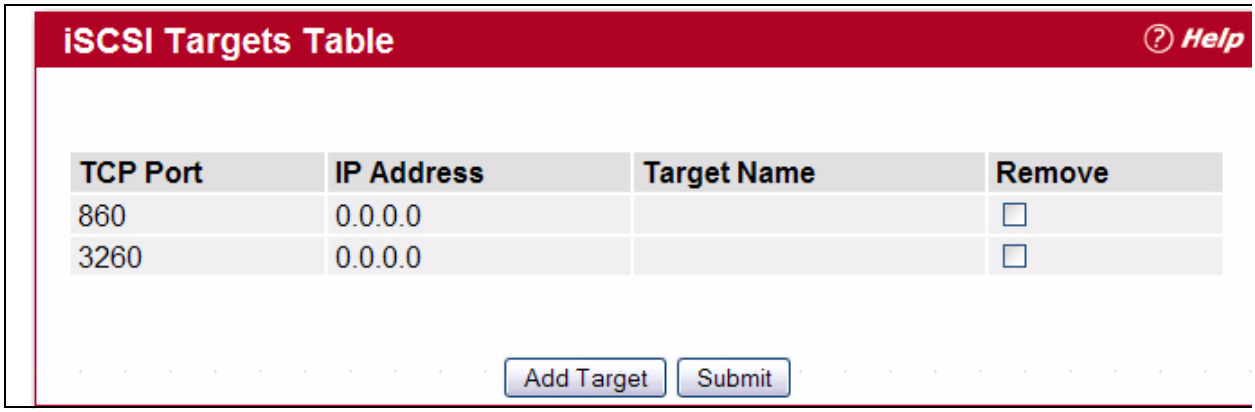

Figure 6-41: iSCSI Targets Table

Table 6-34: iSCSI Targets Table Fields

Field	Description
TCP Port	Shows the TCP port numbers for the targets monitoring iSCSI traffic. The well-known iSCSI ports 3260 and 860 are configured as the default ports. The default ports can be removed. Up to 16 TCP ports can be defined in the system.
IP Address	Shows the IP address of the iSCSI target.
Target Name	Shows the name assigned to the Target.
Remove	Select the checkbox associated with a target and click the **Submit** button to remove a target.
Add Target	Click the **Add Target** button to add an new iSCSI target.

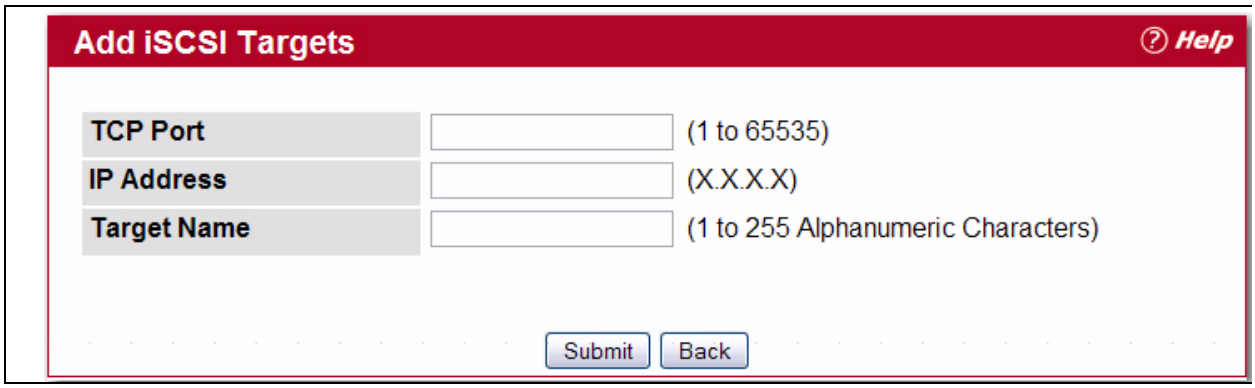

Figure 6-42: Add iSCSI Targets

Table 6-35: Add iSCSI Targets Fields

Field	Description
TCP Port	Enter the TCP port number for the target that will monitor for iSCSI traffic.
IP Address	Enter an IP address the target that will monitor for iSCSI traffic.
Target Name	Enter a name to assign to the Target.

6.5.3 iSCSI Sessions

Use the iSCSI Sessions page to view iSCSI sessions.

To display the iSCSI Sessions page, click **Quality of Service > iSCSI > Sessions** in the navigation menu.

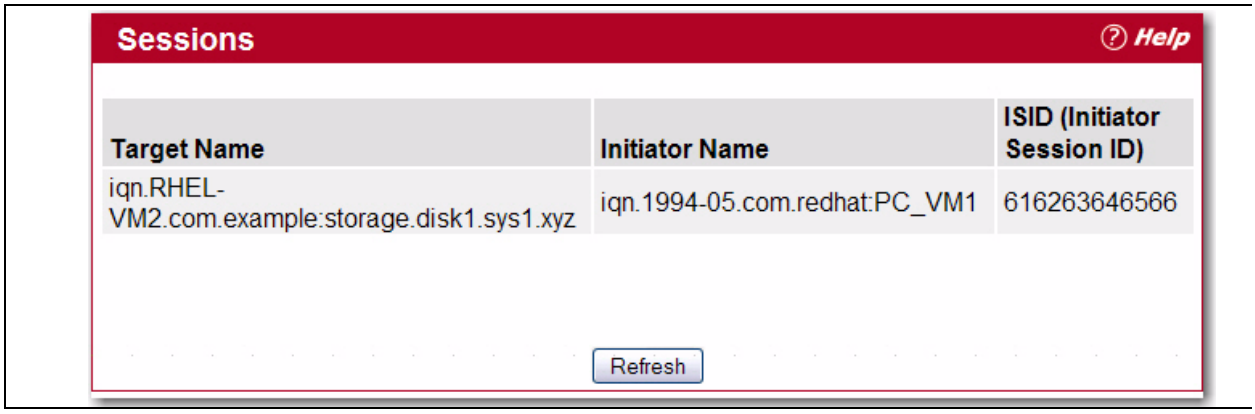

Figure 6-43: iSCSI Sessions

Table 6-36: iSCSI Sessions Fields

Field	Description
Target Name	Shows the name assigned to the Target.
Initiator Name	Shows the name of the initiator.
ISID (Initiator Session ID)	Shows the unique identifier an initiator assigns to the session endpoint. When it is combined with the iSCSI initiator name, it provides a unique name in the world for the SCSI initiator port.

6.5.4 iSCSI Sessions Detailed

Use the iSCSI Sessions Detailed page to view detailed information on iSCSI sessions.

To display the iSCSI Sessions Detailed page, click **Quality of Service > iSCSI > Sessions Detailed** in the navigation menu.

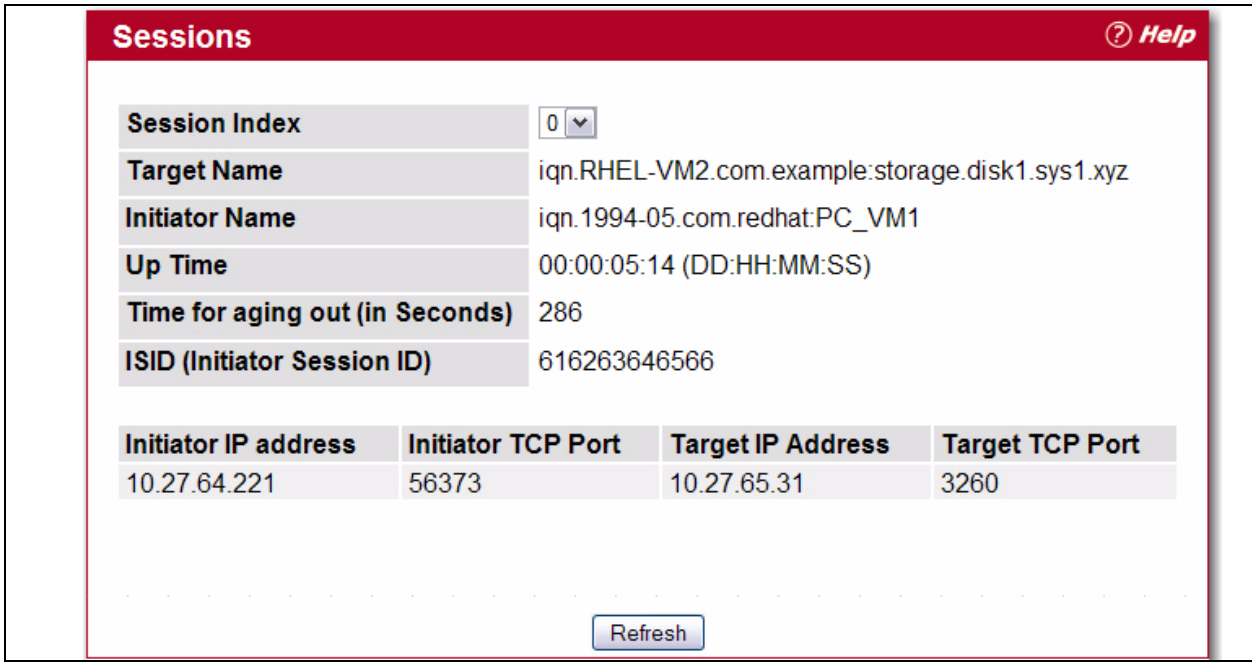

Figure 6-44: iSCSI Sessions Detailed

Table 6-37: iSCSI Sessions Detailed Fields

Field	Description
Session Index	Shows the list of Session indices. The information displayed in the other fields corresponds to the selected Session Index.
Target Name	Shows the name assigned to the Target.
Initiator Name	Shows the name of the initiator.
Up Time	Show time that has elapsed since the session was created.
Time for aging out (in Seconds)	Shows the time (in seconds) left before the session is set to expire.

Table 6-37: iSCSI Sessions Detailed Fields (Continued)

Field	Description
ISID (Initiator Session ID)	Shows the unique identifier an initiator assigns to the session endpoint. When it is combined with the iSCSI initiator name, it provides a unique name in the world for the SCSI initiator port.
Initiator IP Address	Shows the Initiator IP Address.
Initiator TCP Port	Shows the Initiator TCP Port number of one of the connections between the Target and initiator.
Target IP Address	Shows the IP Address of the Target.
Target TCP Port	Shows the Target TCP Port number of one of the connections between the Target and Initiator.

7 Configuring IP Multicast

Multicast protocols are used to deliver Multicast packets from one source to multiple receivers. They facilitate better bandwidth utilization and help to lessen host and router processing, making them ideal for use in application such as video/ audio conferencing, whiteboard tools, stock distribution tickers, etc.

The **IP Multicast** navigation tree menu contains links to the following procedures:

- Configuring DVMRP
- Configuring IGMP
- Managing Multicast Parameters
- Enabling and Configuring PIM-DM
- Enabling and Configuring PIM-SM

Multicast applications send one copy of a packet and address it to a group of receivers (Multicast Group Address) that want to receive it rather than to a single receiver (unicast address). Multicast depends on the network to forward the packets to only those networks and hosts that need to receive them.

Multicast-capable/enabled routers forward multicast packets based on the routes in the Multicast Routing Information Base (MRIB). These routes are created in the MRIB during the process of building multicast distribution trees by the Multicast Protocols running on the router. Different IP Multicast routing protocols use different techniques to construct these multicast distribution trees.

7.1 Managing Multicast Parameters

The multicast protocol maintains a multicast forward table to assist in the forwarding of the multicast packets. The multicast route table contains the forward entries on which the forwarding decision for a stream is made. Each entry in the table refers to a particular multicast stream that is identified by the source IP address and the destination group IP address and each entry is associated with an incoming interface and a set of outgoing interfaces. The routes in the multicast forwarding table are created with the assistance from the multicast routing protocols that take the decision to assign the various outgoing interfaces for a given multicast stream. Multicast routing protocols also determine the incoming interface and the outgoing interface list.

A multicast forwarding cache (MFC) maintains a timer to let multicast routing protocols know about the usage of a specific cache entry being used for forwarding or not. This helps in removing unused mroute entries.

The Multicast protocol also maintains information on network boundaries, or interfaces beyond which multicast messages do not cross.

You can use the Multicast pages to configure parameters that control how the Multicast protocol operates on the network and to view and configure the multicast route (MRoute) table. The Multicast folder provides links to the following pages:

- Multicast Global Configuration
- Multicast Interface Configuration
- Multicast Admin Boundary Configuration
- Multicast Admin Boundary Summary
- Multicast Route Table

- Multicast Static MRoute Configuration
- Multicast Static MRoute Table SummaryEnabling and Configuring PIM-DM

7.1.1 Multicast Global Configuration

Use this page to enable and disable multicast operation across the router and to view summary statistics about multicast operation. To display this page, click **IPv4 Multicast > Global Configuration** in the navigation tree.

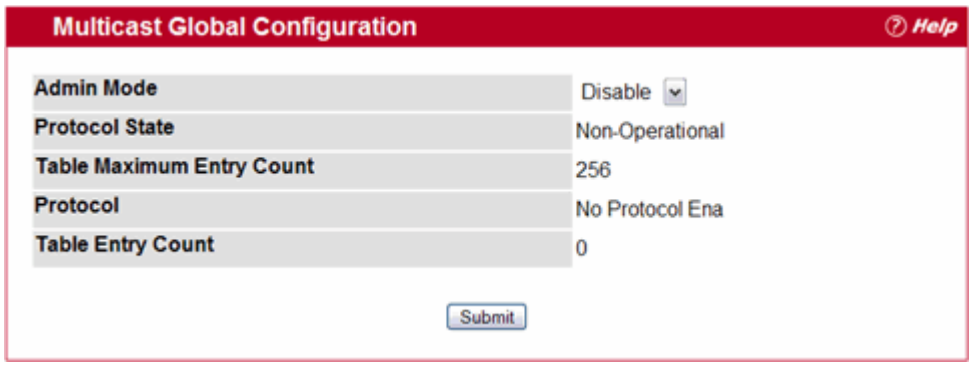

Figure 7-1: Multicast Global Configuration

Table 7-1: Multicast Global Configuration Fields

Field	Description
Admin Mode	Select enable or disable to set the administrative status of Multicast Forwarding in the router. The default is disable.
Protocol State	The operational state of the multicast forwarding module.
Table Maximum Entry Count	The maximum number of entries in the IP Multicast routing table.
Protocol	The multicast routing protocol presently activated on the router, if any.
Table Entry Count	The number of multicast route entries currently present in the Multicast route table.

Click **Submit** to send the updated configuration to the switch. Configuration changes take effect immediately. These changes will not be retained across a power cycle unless a save is performed.

7.1.2 Multicast Interface Configuration

Use this page to enable and disable multicast operation on specific ports. To display this page, click **IPv4 Multicast > Interface Configuration** in the navigation tree.

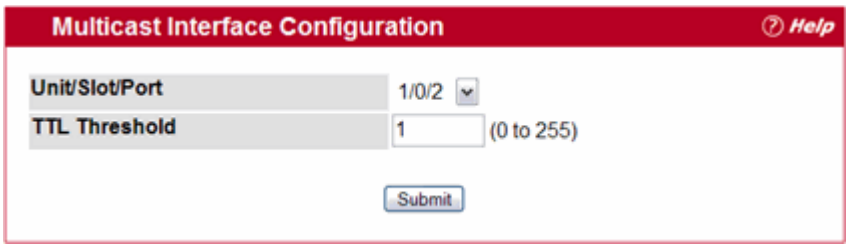

Figure 7-2: Multicast Interface Configuration

Table 7-2: Multicast Interface Configuration

Field	Description
Slot/Port	Select the routing interface you want to configure from the menu.
TTL Threshold	Enter the TTL threshold below which a multicast data packet will not be forwarded from the selected interface. You should enter a number between 0 and 255. If you enter 0 all multicast packets for the selected interface will be forwarded. You must configure at least one router interface before you will see this field.

Click **Submit** to send the updated configuration to the switch. Configuration changes take effect immediately. These changes will not be retained across a power cycle unless a save is performed.

7.1.3 Multicast Admin Boundary Configuration

You can configure a router interface as a multicast boundary to stop the ingress and egress of multicast traffic for a given range of multicast addresses over that interface. To display this page, click **IPv4 Multicast > Admin Boundary Configuration** in the navigation tree.

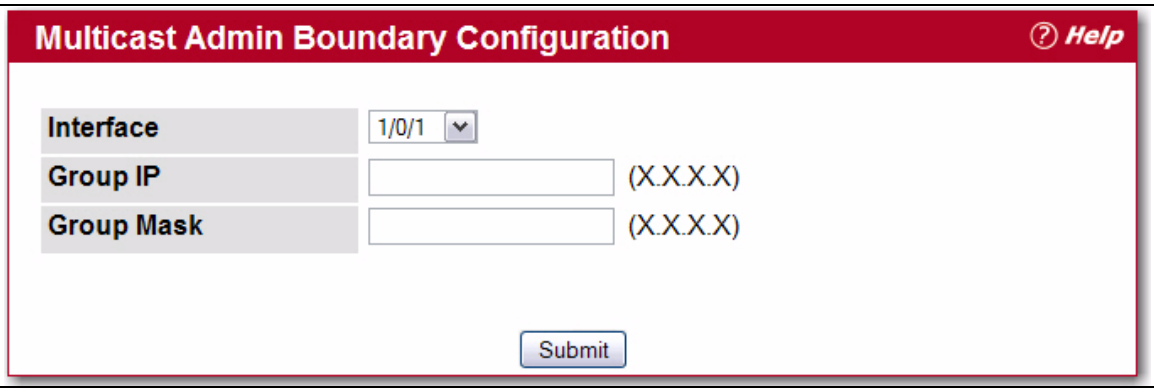

Figure 7-3: Multicast Admin Boundary Configuration

Table 7-3: Multicast Admin Boundary Configuration Fields

Field	Description
Interface	Select the router interface for which the administratively scoped boundary is to be configured.
Group IP	Enter the multicast group address for the start of the range of addresses to be excluded. The address must be in the range of 239.0.0.0 through 239.255.255.255.
Group Mask	Enter the mask to be applied to the multicast group address. The combination of the mask and the Group IP gives the range of administratively scoped addresses for the selected interface.

- Click **Submit** to send the updated configuration to the switch. Configuration changes take effect immediately. These changes will not be retained across a power cycle unless a save is performed.
- To delete the selected administrative scoped boundary, click **Delete**.

7.1.4 Multicast Admin Boundary Summary

This page displays each interface that is configured as a multicast boundary and lists the multicast groups for which the interface serves as a boundary. To display this page, click **IPv4 Multicast > Admin Boundary Summary** in the navigation tree.

Figure 7-4: Multicast Admin Boundary Summary

Table 7-4: Multicast Admin Boundary Summary Fields

Field	Description
Interface	The router interface to which the administratively scoped address range is applied.
Group IP	The multicast group address for the start of the range of addresses to be excluded.
Group Mask	The mask that is applied to the multicast group address. The combination of the mask and the Group IP gives the range of administratively scoped addresses for the selected interface.
Delete	Deletes the selected admin boundary scope entries in the router.
Refresh	Refresh the data on the screen with the present state of the data in the router.

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

7.1.5 Multicast Route Table

This screen displays selected contents of the Mroute table in tabular form. If there are no routes in the table you will not be presented with the selection criteria. To display this page, click **IPv4 Multicast > Multicast Route Table** in the navigation tree.

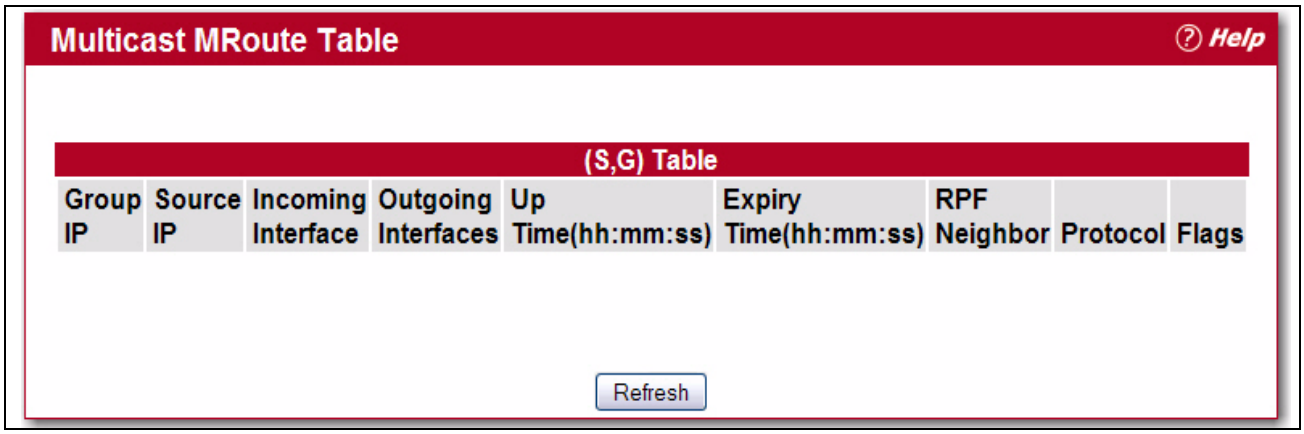

Figure 7-5: Multicast Route Table

Table 7-5: Multicast Route Table Fields

Field	Description
Source IP	Enter the IP address of the multicast packet source to be combined with the Group IP to fully identify a single route whose Mroute table entry you want to display or clear. You may leave this field blank.
Group IP	Enter the destination group IP address whose multicast route(s) you want to display or clear.
Incoming Interface	The incoming interface on which multicast packets for this source/group arrive.
Outgoing Interface(s)	The list of outgoing interfaces on which multicast packets for this source/group are forwarded.
Up Time	The time in seconds since the entry was created.
Expiry Time	The time in seconds before this entry will age out and be removed from the table.
RPF Neighbor	The IP address of the Reverse Path Forwarding neighbor.
Protocol	The multicast routing protocol which created this entry. The possibilities are: • PIM-DM • PIM-SM • DVMRP
Flags	The value displayed in this field is valid if the multicast routing protocol running is PIMSM. The possible values are RPT or SPT. For other protocols an "------" is displayed.

- Click **Search** to search the Mroute table for an entry matching the Source IP (if entered) and Group IP address.
- Click **Clear Route** to remove the data on the screen for the Source IP (if entered) and Group IP address you have specified.
- Click **Clear All** to remove all the data on the screen.
- Click **Refresh** to refresh the information on the screen with the present state of the data in the router.

7.1.6 Multicast Static MRoute Configuration

You can use this page to add static multicast routes to the MRoute table. To display this page, click **IPv4 Multicast > Static MRoute Table Configuration** in the navigation tree.

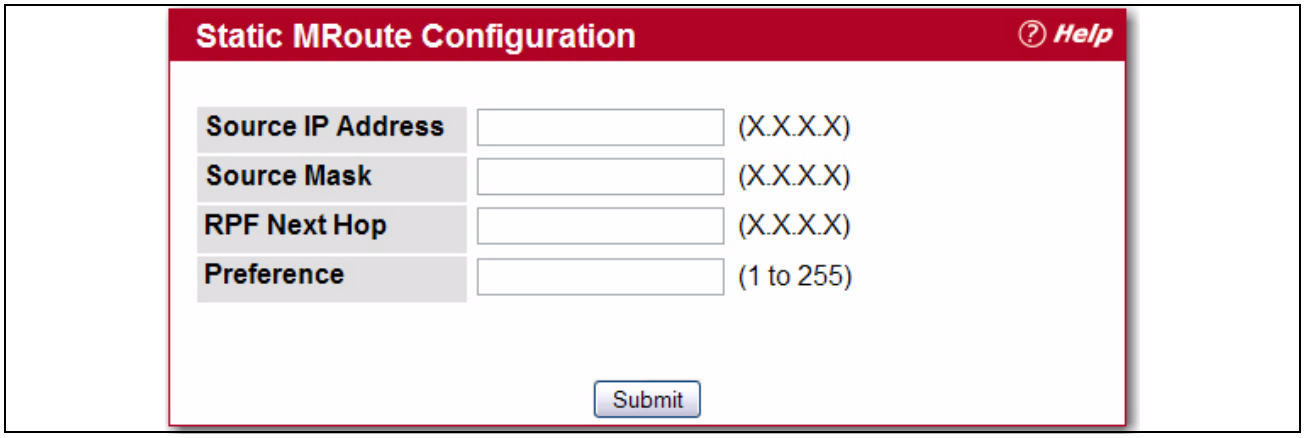

Figure 7-6: Static MRoute Configuration

Table 7-6: Static MRoute Configuration Fields

Field	Description
Source IP Address	The address of the Multicast data source.
Source Mask	The network mask for the IP address of the Multicast data source to be configured.
RPF Next Hop	The RPF Address for the source range of static mroute entry.
Preference	The preference with which the static mroute to be considered against other matching static mroute entry for a given source.

Click **Submit** to send the updated configuration to the switch. Configuration changes take effect immediately. These changes will not be retained across a power cycle unless a save is performed.

7.1.7 Multicast Static MRoute Table Summary

This page lists all static routes. To display it, click **IPv4 Multicast > Static MRoute Table Summary** in the navigation tree.

Figure 7-7: Static MRoute Table Summary

Table 7-7: Static MRoute Table Summary Fields

Field	Description
Source IP	The address of the Multicast data source.
Source Mask	The network mask for the IP address of the Multicast data source to be configured.
RPF Next Hop	The RPF Address for the source range of static mroute entry.
Preference	The preference with which the static mroute to be considered against other matching static mroute entry for a given source.

- To delete the selected static mroute entries in the router, click **Delete**.
- Click **Refresh** to update the information on the screen.

7.2 Configuring DVMRP

Note...

The DVMRP protocol can be configured to operate on IPv4 networks only.

The Distance Vector Multicast Routing Protocol (DVMRP) is a distributed protocol. The DVMRP protocol operates as follows:

- The first message for any source-group pair is forwarded to the entire multicast network, with respect to the time-to-live (TTL) of the packet.
- TTL restricts the area to be flooded by the message.
- All the leaf routers that do not have members on directly attached subnetworks send back prune messages to the upstream router.
- The branch that transmitted a prune message is deleted from the delivery tree.
- The delivery tree, which is spanning to all the members in the multicast group, is constructed.

The DVMRP folder contains links to the following pages:

- DVMRP Global Configuration
- Interface Configuration
- DVMRP Configuration Summary
- Next Hop Summary
- Prune Summary
- Route Summary

7.2.1 DVMRP Global Configuration

Use the Global Configuration page to administratively enable and disable the feature and to view basic information on its operation in the network.

To access the page, click **IPv4 Multicast > DVMRP > Global Configuration** in the navigation tree.

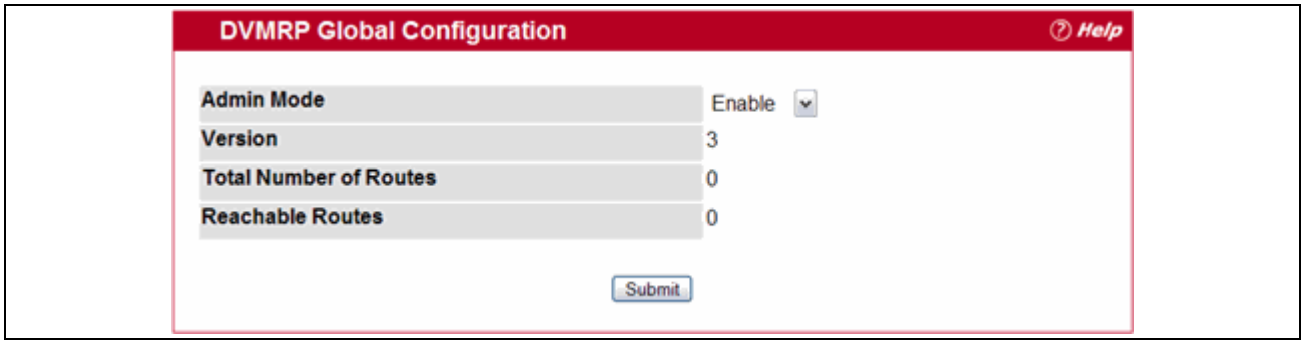

Figure 7-8: DVMRP Global Configuration

Table 7-8: DVMRP Global Configuration Fields

Field	Description
Admin Mode	Sets the administrative status of DVMRP to active (**Enable**) or inactive (**Disable**). The default is Disable.
Version	Displays the DVMRP version.
Total Number of Routes	Displays the number of routes in the DVMRP routing table.
Reachable Routes	Displays the number of routes in the DVMRP routing table that have a non-infinite metric (i.e., they are reachable).

If you change the administrative mode, click **Submit** to send the updated configuration to the switch. Configuration changes take effect immediately. These changes will not be retained across a power cycle unless a save is performed.

7.2.2 Interface Configuration

You can enable a DVMRP interface on any configured IP interface. You must configure at least one router interface before you configure a DVMRP interface.

To access the page, click **IPv4 Multicast > DVMRP > Interface Configuration** in the navigation tree.

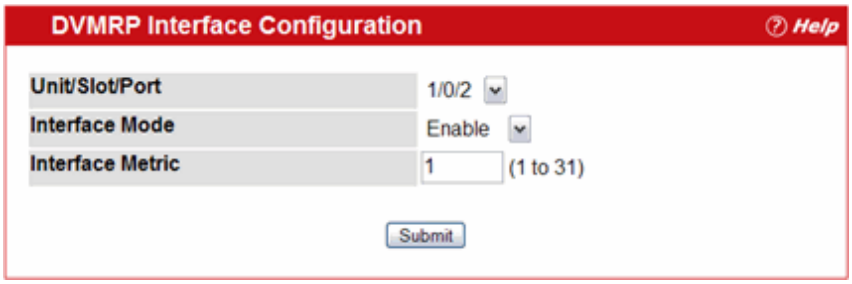

Figure 7-9: DVMRP Interface Configuration

Table 7-9: DVMRP Interface Configuration Fields

Field	Description
Slot/Port	Selects the interface for which data is to be configured. If no router interfaces have been set up, a message displays that no router interfaces are available, and the configuration screen will not be displayed.
Interface Mode	Enables or disables the administrative mode of the selected DVMRP routing interface.
Interface Metric	Specifies the DVMRP metric for the selected interface. This value is sent in DVMRP messages as the cost to reach this network. Valid values are from 1 to 31.

If you change the administrative mode, click **Submit** to send the updated configuration to the switch. Configuration changes take effect immediately. These changes will not be retained across a power cycle unless a save is performed.

7.2.3 DVMRP Configuration Summary

The Configuration Summary page displays information on a selected DVMRP-enabled interface. You must configure at least one router interface before you can display data for a DVMRP interface. Otherwise a message displays that no router interfaces are available, and the configuration summary screen is not displayed.

To display this page, click **IPv4 Multicast > DVMRP > Configuration Summary** in the navigation tree.

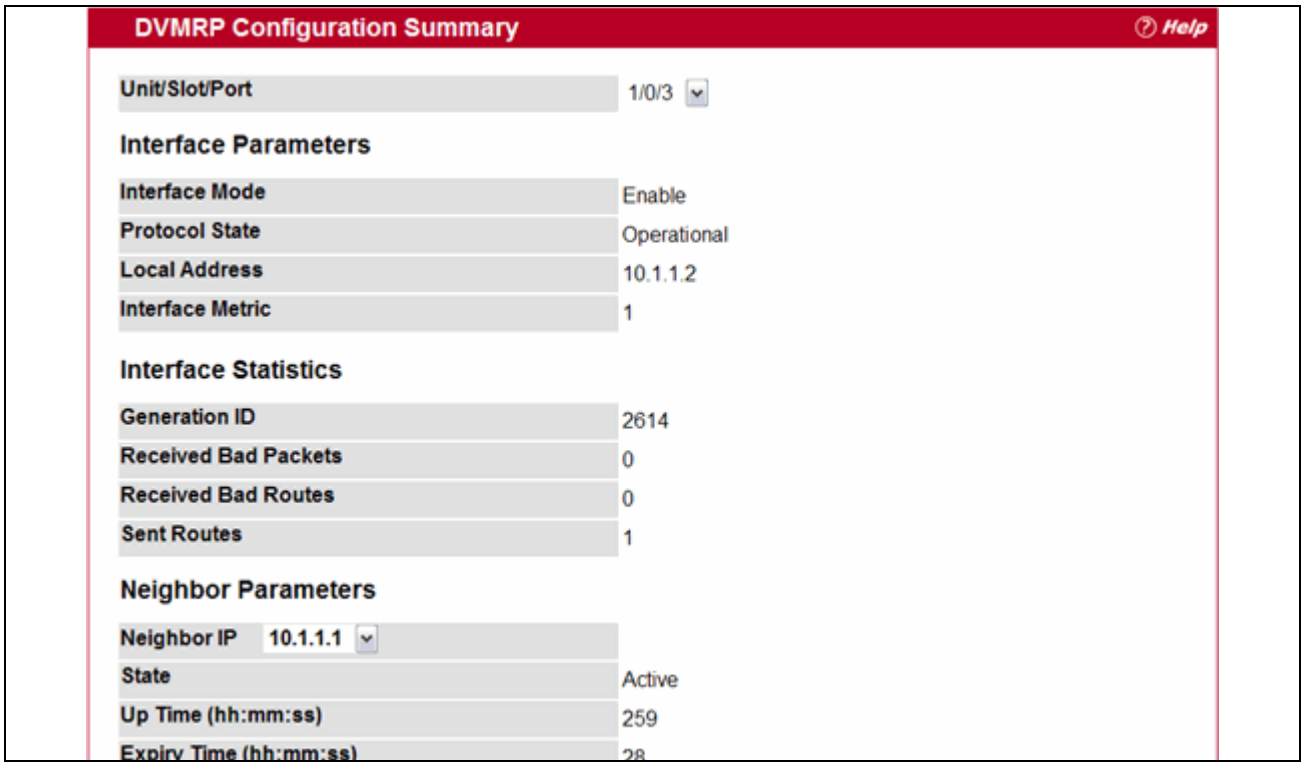

Figure 7-10: DVMRP Configuration Summary

Table 7-10: DVMRP Configuration Summary Fields

Field	Description
Slot/Port	Selects the interface for which data is to be displayed.
Interface Mode	The administrative mode (Enabled or Disabled) of the selected DVMRP routing interface.
Protocol State	The operational state of the DVMRP protocol on the selected interface, either operational or non-operational.
Local Address	The IP address used as a source address in DVMRP packets sent from the selected interface.
Interface Metric	The metric used to calculate distance vectors for the selected interface.
Generation ID	The DVMRP generation ID used by the router for the selected interface. This value is reset every time an interface is (re)started and is placed in prune messages. A change in generation ID informs the neighbor routers that any previous information about this router should be discarded.
Received Bad Packets	The number of invalid packets received on the selected interface.
Received Bad Routes	The number of invalid routes received on the selected interface.
Sent Routes	The number of routes sent on the selected interface.
Neighbor IP	The IP address of the neighbor whose information is displayed.
State	The state of the specified neighbor router on the selected interface, either active or down.
Neighbor Uptime	The DVMRP uptime for the specified neighbor on the selected interface. This is the time since the neighbor entry was learned.
Neighbor Expiry Time	The DVMRP expiry time for the specified neighbor on the selected interface. This is the time left before this neighbor entry will age out, and is not applicable if the neighbor router's state is down.
Generation ID	The DVMRP generation ID for the specified neighbor on the selected interface.
Major Version	The DVMRP Major Version for the specified neighbor on the selected interface.

Table 7-10: DVMRP Configuration Summary Fields (Continued)

Field	Description
Minor Version	The DVMRP Minor Version for the specified neighbor on the selected interface.
Capabilities	The DVMRP capabilities of the specified neighbor on the selected interface.
Received Routes	The number of routes received for the specified neighbor on the selected interface.
Received Bad Packets	The number of invalid packets received for the specified neighbor on the selected interface.
Received Bad Routes	The number of invalid routes received for the specified neighbor on the selected interface.

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

7.2.4 Next Hop Summary

The DVMRP Next Hop Summary page lists each source IP address/network for which the router maintains a DVMRP route, and the next hop for DVMRP messages originating from that source. To display this page, click **IPv4 Multicast > DVMRP > Next Hop Summary** in the navigation tree.

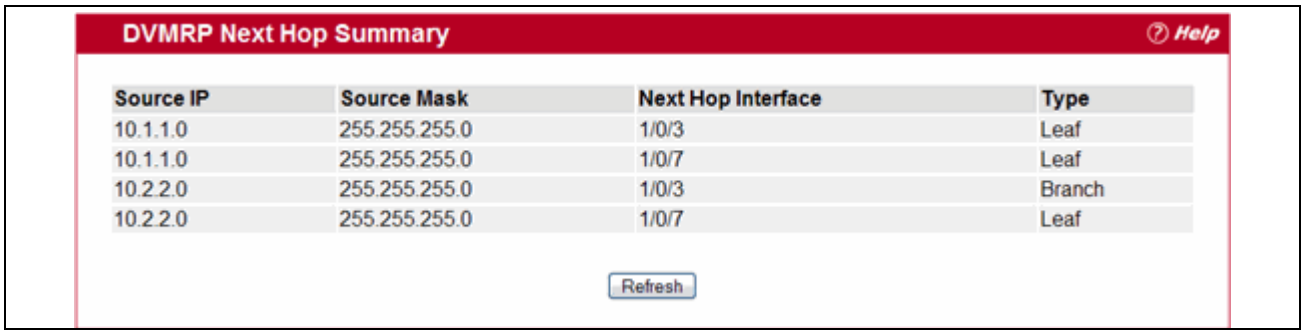

Figure 7-11: DVMRP Next Hop Summary

Table 7-11: DVMRP Next Hop Summary Fields

Field	Description
Source IP	The IP address of the source network for this table entry.
Source Mask	The network mask for the Source IP address.
Next Hop Interface	The outgoing interface for this next hop.
Type	The next hop type. “Leaf” means that no downstream-dependent neighbors exist on the outgoing interface. Otherwise, the type is “branch”

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

7.2.5 Prune Summary

Edge routers remove their multicast interfaces that do not have any members for a particular multicast group. When all downstream interfaces are removed, the router informs its upstream router, asking it not to send any more traffic destined to that network. The message used to convey this is called a Prune Message. The upstream router removes this interface from its downstream interface list if it receives prune messages from each downstream router. If a router removes all downstream interfaces, it can send prune to its upstream router.

Every prune message has lifetime, after which, the interface is joined back onto delivery tree. If the unwanted datagrams still appear, the prune is reinitiated. In a scenario where all the interfaces are pruned and a prune is sent upstream, its lifetime will be minimum of all the lifetimes of prunes of downstream interfaces.

The DVMRP Prune Summary page provides information for each active prune message. To display this page, click **IPv4 Multicast > DVMRP > Prune Summary** in the navigation tree.

Figure 7-12: DVMRP Prune Summary

Table 7-12: DVMRP Prune Summary Fields

Field	Description
Group IP	The group address which has been pruned.
Source IP	The address of the source or source network which has been pruned.
Source Mask	The subnet mask to be combined with the source IP address to identify the source or source network which has been pruned.
Expiry Time	The amount of time remaining before this prune should expire at the upstream neighbor. If no prune messages have been received from downstream neighbors, this is set to the value of the default prune lifetime timer; otherwise, it is set to the smallest received value or the default timer, whichever is less.

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

7.2.6 Route Summary

DVMRP-enabled interfaces discover their neighbors by periodically sending probe messages. These probe messages contain the list of its neighboring DVMRP routers from which a probe has been received. Thus, a DVMRP routing table is created.

The DVMRP Route Summary page lists data for each DVMRP route. To display this page, click **IPv4 Multicast > DVMRP > Route Summary** in the navigation tree.

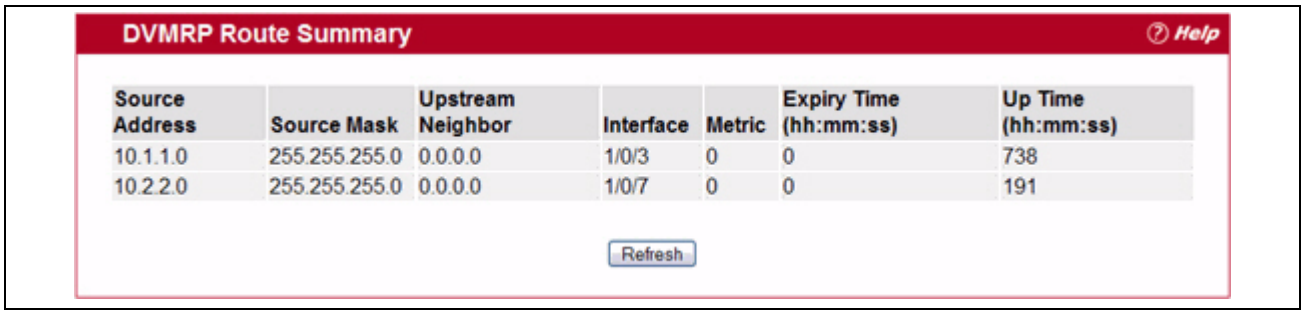

Source Address	Source Mask	Upstream Neighbor	Interface	Metric	Expiry Time (hh:mm:ss)	Up Time (hh:mm:ss)
10.1.1.0	255.255.255.0	0.0.0.0	1/0/3	0	0	738
10.2.2.0	255.255.255.0	0.0.0.0	1/0/7	0	0	191

Figure 7-13: DVMRP Route Summary

Table 7-13: DVMRP Route Summary Fields

Field	Description
Source Address	The network address that is combined with the source mask to identify the sources for this entry.
Source Mask	The subnet mask to be combined with the source address to identify the sources for this entry.
Upstream Neighbor	The address of the upstream neighbor (e.g., RPF neighbor) from which IP datagrams from these sources are received.
Interface	The interface on which IP datagrams sent by these sources are received. A value of 0 typically means the route is an aggregate for which no next-hop interface exists.
Metric	The distance in hops to the source subnet.
Expiry Time	The minimum amount of time remaining before this entry will be aged out.
Up Time	The time since the route represented by this entry was learned by the router.

7.3 Configuring IGMP

IGMP enables a multicast router to learn which multicast addresses are of interest to the systems connected to each of its directly-attached networks. IGMP version 3 also adds the capability for a multicast router to learn which sources are of interest to neighboring systems for packets sent to any particular multicast address. The information gathered by IGMP is provided to the multicast routing protocol (i.e., DVMRP, PIM-DM, and PIM-SM) that is currently active on the router in order to ensure multicast packets are delivered to all networks where there are interested receivers.

The IGMP folder contains links to the following pages:

- IGMP Global Configuration
- IGMP Routing Interface Configuration
- IGMP Proxy Configuration

7.3.1 IGMP Global Configuration

Use this page to administratively turn on and off IGMP processing on the router. To display this page, click **IPv4 Multicast > IGMP > Routing Interface > Global Configuration** in the navigation tree.

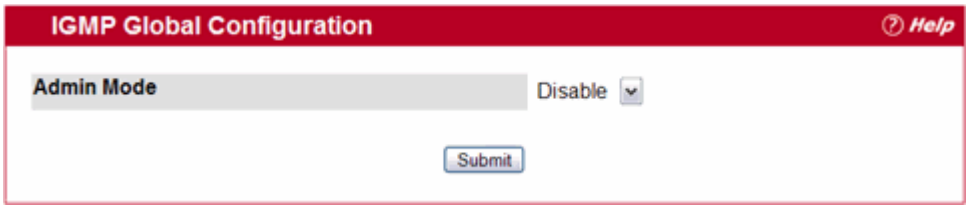

Figure 7-14: IGMP Global Configuration

Table 7-14: IGMP Global Configuration Fields

Field	Description
Admin Mode	Select enable or disable from the pulldown menu to set the administrative status of IGMP in the router to active or inactive. The default is disable.

Click **Submit** to send the updated configuration to the switch. Configuration changes take effect immediately. These changes will not be retained across a power cycle unless a save is performed.

7.3.2 IGMP Routing Interface Configuration

The IGMP Routing Interface Configuration folder contains links to the following features:

- IGMP Interface Configuration
- IGMP Configuration Summary
- IGMP Cache Information
- Configuration Summary

7.3.2.1 IGMP Interface Configuration

In order for subnets and hosts on a network to participate in multicast groups, they must be connected to an IGMP-enabled router interface. Use this page to enable IGMP on router interfaces and to configure properties of that interface.

To display this page, click **IPv4 Multicast > IGMP > Routing Interface > Interface Configuration** in the navigation tree.

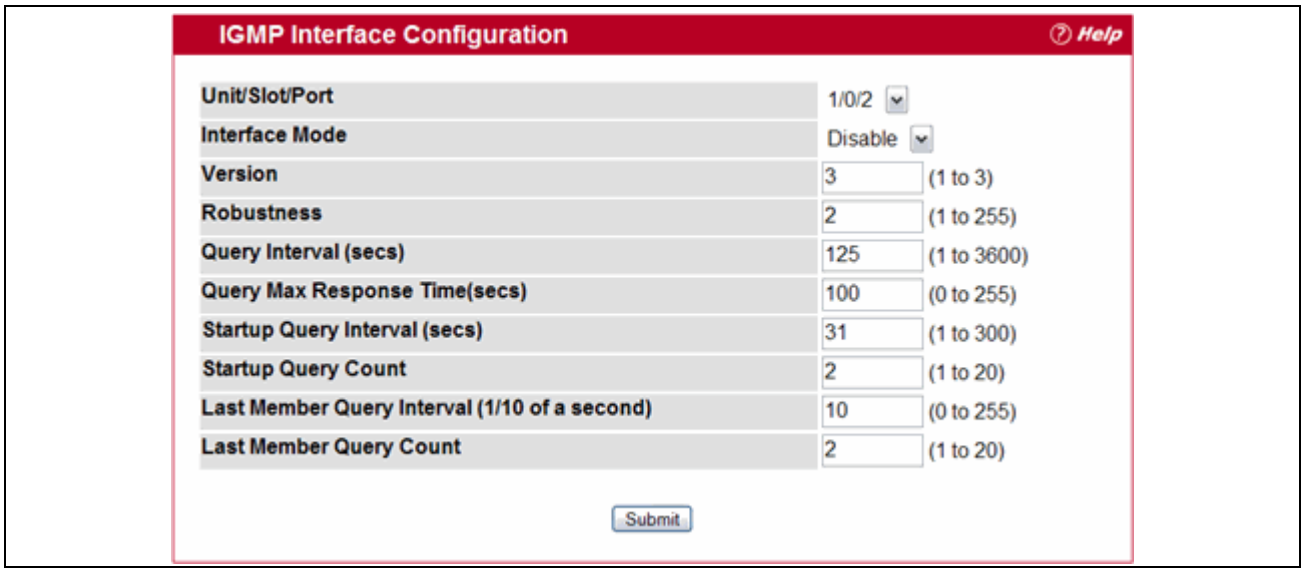

Figure 7-15: IGMP Interface Configuration

Table 7-15: IGMP Interface Configuration Fields

Field	Description
Slot/Port	Select the slot and port for which data is to be displayed or configured from the pulldown menu. Slot 0 is the base unit. You must have configured at least one router interface before configuring or displaying data for an IGMP interface, otherwise an error message will be displayed.
Interface Mode	Select enable or disable from the pulldown menu to set the administrative status of IGMP on the selected interface. The default is disable.
Version	Enter the version of IGMP you want to configure on the selected interface. Valid values are (1 to 3) and the default value is 3. This field is configurable only when IGMP interface mode is enabled.
Robustness	Enter the robustness value. This variable allows tuning for the expected packet loss on a subnet. If you expect the subnet to be lossy, you should enter a higher number for this parameter. IGMP is robust to (robustness variable-1) packet losses. Valid values are from 1 to 255. The default value is 2.
Query Interval	Enter the frequency in seconds at which IGMP host-query packets are to be transmitted on this interface. Valid values are from 1 to 3600. The default value is 125.
Query Max Response Time	Enter the maximum query response time to be advertised in IGMPv2 queries on this interface, in tenths of a second. The default value is 10. Valid values are from (0 to 255).
Startup Query Interval	Enter the number of seconds between the transmission of startup queries on the selected interface. The valid values are from 1 to 300. The default value is 31.
Startup Query Count	Enter the number of queries to be sent on startup. The valid values are from 1 to 20. The default value is 2.
Last Member Query Interval	Enter the last member query interval in tenths of a second. This the maximum response time to be inserted into group-specific queries sent in response to leave group messages, and is also the amount of time between group-specific query messages. Valid values are from 0 to 255. The default value is 1. This value is not used for IGMP version 1.
Last Member Query Count	Enter the number of queries to be sent on receiving a leave group report. Valid values are from 1 to 20. The default value is 2.

Click **Submit** to send the updated configuration to the switch. Configuration changes take effect immediately. These changes will not be retained across a power cycle unless a save is performed.

7.3.2.2 IGMP Configuration Summary

The IGMP Configuration Summary page displays configuration information for the selected IGMP-enabled router interface.

To display this page, click **IPv4 Multicast > IGMP > Routing Interface > Configuration Summary** in the navigation tree.

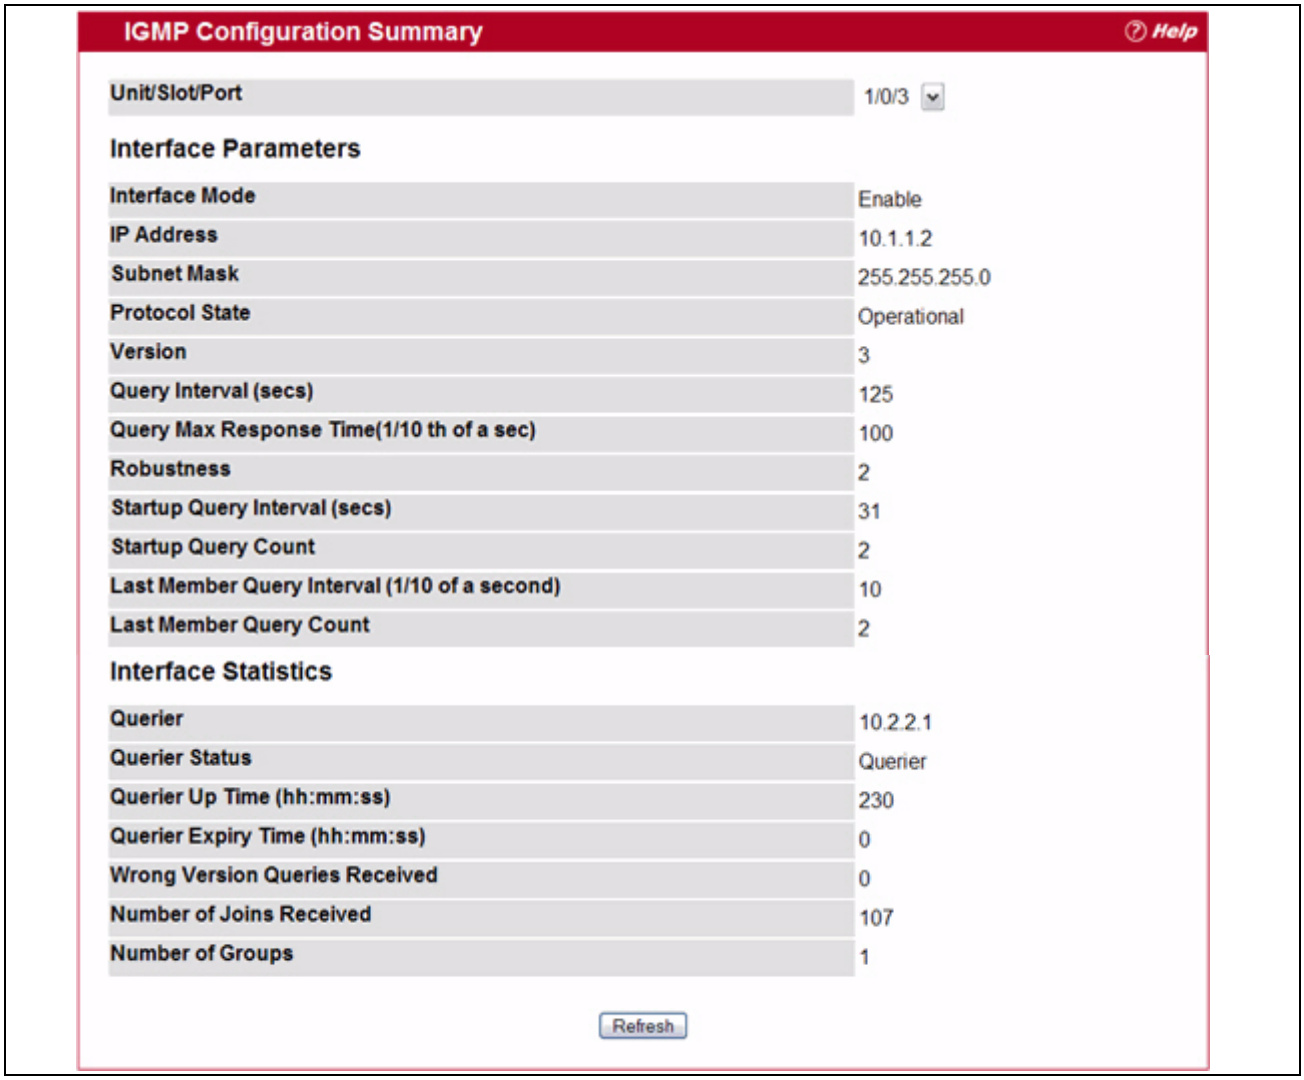

Figure 7-16: IGMP Configuration Summary

Table 7-16: IGMP Configuration Summary Fields

Field	Description
Slot/Port	Select the slot and port for which data is to be displayed. Slot 0 is the base unit.
Interface Mode	The administrative status of IGMP on the selected interface.
IP Address	The IP address of the selected interface.
Subnet Mask	The subnet mask for the IP address of the selected interface.
Protocol State	The operational state of IGMP on the selected interface.
Version	The version of IGMP configured on the selected interface.

Table 7-16: IGMP Configuration Summary Fields (Continued)

Field	Description
Query Interval	The frequency at which IGMP host-query packets are transmitted on the selected interface.
Query Max Response Time	The maximum query response time advertised in IGMPv2 queries sent from the selected interface.
Robustness	The robustness parameter for the selected interface. This variable allows tuning for the expected packet loss on a subnet. If a subnet is expected to be lossy, the robustness variable may be increased. IGMP is robust to (robustness variable-1) packet losses.
Startup Query Interval	The interval at which startup queries are sent on the selected interface.
Startup Query Count	The number of queries to be sent on startup.
Last Member Query Interval	The last member query interval. The last member query interval is the maximum response time inserted into group-specific queries sent in response to leave group messages, and is also the amount of time between group-specific query messages. This value may be tuned to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group. This value is not used for IGMP version 1.
Last Member Query Count	The number of queries to be sent on receiving a leave group report.
Querier	The address of the IGMP querier on the IP subnet to which the selected interface is attached.
Querier Status	Indicates whether the selected interface is in querier or non querier mode.
Querier Up Time	The time in seconds since the IGMP interface querier was last changed.
Querier Expiry Time	The time in seconds remaining before the other querier present timer expires. If the local system is the querier, this will be zero.
Wrong Version Queries	The number of queries that have been received on the selected interface with an IGMP version that does not match the IGMP version configured for the interface, over the lifetime of the entry. IGMP requires that all routers on a LAN be configured to run the same version of IGMP. Therefore, a configuration error is indicated if any queries are received with the wrong version number.
Number of Joins	The number of times a group membership has been added on the selected interface; that is, the number of times an entry for this interface has been added to the cache table. This gives an indication of the amount of IGMP activity on the interface.
Number of Groups	The current number of entries for the selected interface in the cache table.

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

7.3.2.3 IGMP Cache Information

The IGMP router cache table provides information on the IP multicast groups for which there are members on a particular router interface. The entries are learned when the IGMP interface receives a join message from a client, and are removed from the cache after the configured Expiry Time has elapsed.

To display this page, click **IPv4 Multicast > IGMP > Routing Interface > Cache Information** in the navigation tree.

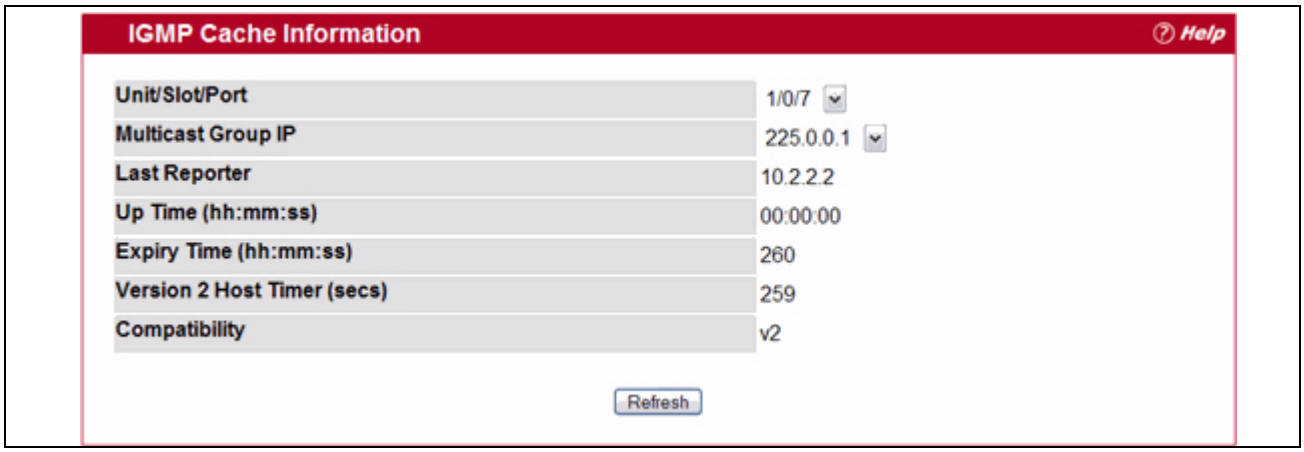

Figure 7-17: IGMP Cache Information

Table 7-17: IGMP Cache Information Fields

Field	Description
Slot/Port	Select the unit, slot, and port for which data is to be displayed. Slot 0 is the base unit.
Multicast Group IP	Select the IP multicast group address for which data is to be displayed. If no group membership reports have been received on the selected interface you will not be able to make this selection, and none of the non-configurable data will be displayed.
Last Reporter	The IP address of the source of the last membership report received for the IP Multicast group address on the selected interface.
Up Time	The time elapsed since this entry was created.
Expiry Time	The minimum amount of time remaining before this entry will be aged out.
Version 1 Host Timer	The time remaining until the local router will assume that there are no longer any IGMP version 1 members on the IP subnet attached to this interface. When an IGMPv1 membership report is received, this timer is reset to the group membership timer. While this timer is non-zero, the local router ignores any IGMPv2 leave messages for this group that it receives on the selected interface. This field is displayed only if the interface is configured for IGMP version 1.
Version 2 Host Timer	The time remaining until the local router will assume that there are no longer any IGMP version 2 members on the IP subnet attached to this interface. When an IGMPv2 membership report is received, this timer is reset to the group membership timer. While this timer is non-zero, the local router ignores any IGMPv1 and IGMPv3 leave messages for this group that it receives on the selected interface. This field is displayed only if the interface is configured for IGMP version 2.
Compatibility	Shows group compatibility mode (v1, v2 and v3) for this group on the specified interface.
Filter Mode	The source filter mode (Include/Exclude/NA) for the specified group on this interface. When NA mode is active the field is blank

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

7.3.2.4 IGMP Interface Source List Information

Use this page to display information on multicast groups associated with an interface and the hosts that belong to each group.

To display this page, click **IPv4 Multicast > IGMP > Routing Interface > Source List Information** in the navigation tree.

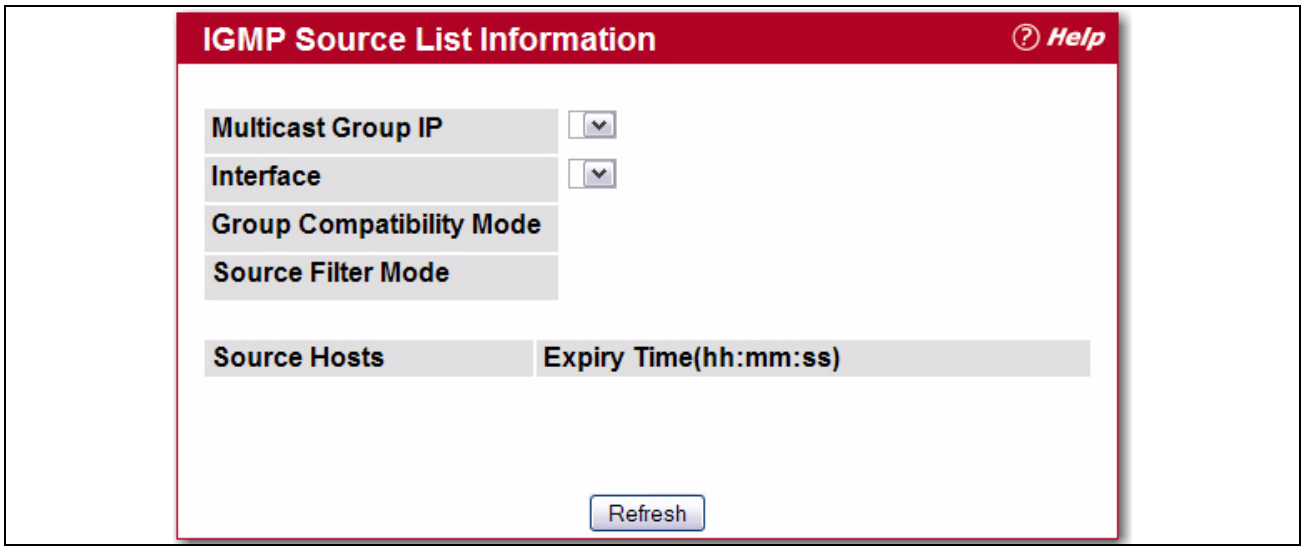

Figure 7-18: IGMP Interface Source List Information

Table 7-18: IGMP Interface Source List Information Fields

Field	Description
Multicast Group IP	Select the IP multicast group address for which data is to be displayed. If no group membership reports have been received on the selected interface you will not be able to make this selection, and none of the non-configurable data will be displayed.
Interface	This parameter shows the interface on which multicast packets are forwarded.
Group Compatibility Mode	This parameter shows group compatibility mode (v1, v2, and v3) for this group on the specified interface.
Source Filter Mode	The source filter mode (Include/Exclude/NA) for the specified group on this interface.
Source Hosts	This parameter shows source addresses which are members of this multicast address.
Expiry Time	This parameter shows expiry time interval against each source address which are members of this multicast group. This is the amount of time after which the specified source entry is aged out.

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

7.3.3 IGMP Proxy Configuration

When you configure an interface in IGMP proxy mode, it acts as a proxy multicast host that sends IGMP membership reports on one interface for IGMP Membership reports received on all other IGMP-enabled router interfaces.

The Proxy Interface folder provides links to the following pages:

- Interface Configuration
- Configuration Summary
- Interface Membership Information
- Interface Membership Information—Detailed

7.3.3.1 Interface Configuration

Use this page to enable and disable ports as IGMP proxy interfaces. To display this page, click **IPv4 Multicast > IGMP > Proxy Interface > Interface Configuration** in the navigation tree.

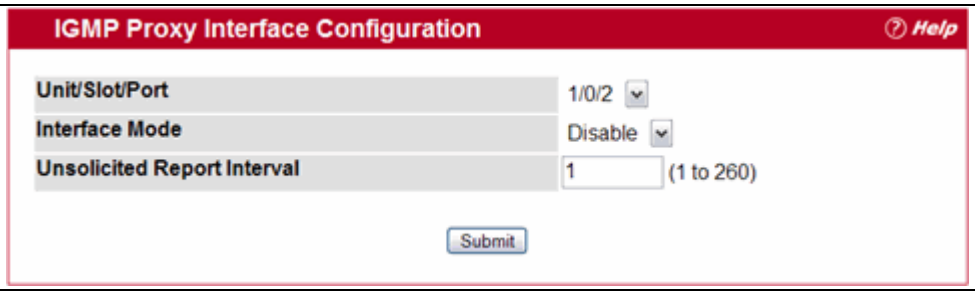

Figure 7-19: IGMP Proxy Interface Configuration

Table 7-19: IGMP Proxy Interface Configuration Fields

Field	Description
Slot/Port	Select the port for which data is to be displayed or configured from the pulldown menu. You must have configured at least one router interface before configuring or displaying data for an IGMP Proxy interface and it should not be a IGMP routing interface. This field is configurable only when interface mode is disabled.
Interface Mode	Select enable or disable from the pulldown menu to set the administrative status of IGMP Proxy on the selected interface. The default is disable. Routing, IGMP and Multicast global admin modes should be enabled to enable IGMP Proxy interface mode.
Version	Enter the version of IGMP you want to configure on the selected interface. Valid values are (1 to 3) and the default value is 3. This field is configurable only when IGMP Proxy interface mode is enabled.
Unsolicited Report Interval	Enter the unsolicited time interval value in seconds. The Unsolicited Report Interval is the time between repetitions of a host's initial report of membership in a group. Valid values are from (1 to 260). The default value is 1.

Click **Submit** to send the updated configuration to the switch. Configuration changes take effect immediately. These changes will not be retained across a power cycle unless a save is performed.

7.3.3.2 Configuration Summary

Use this page to view configuration and statistics on IGMP proxy-enabled interfaces. To display this page, click **IPv4 Multicast > IGMP > Proxy Interface > Configuration Summary** in the navigation tree.

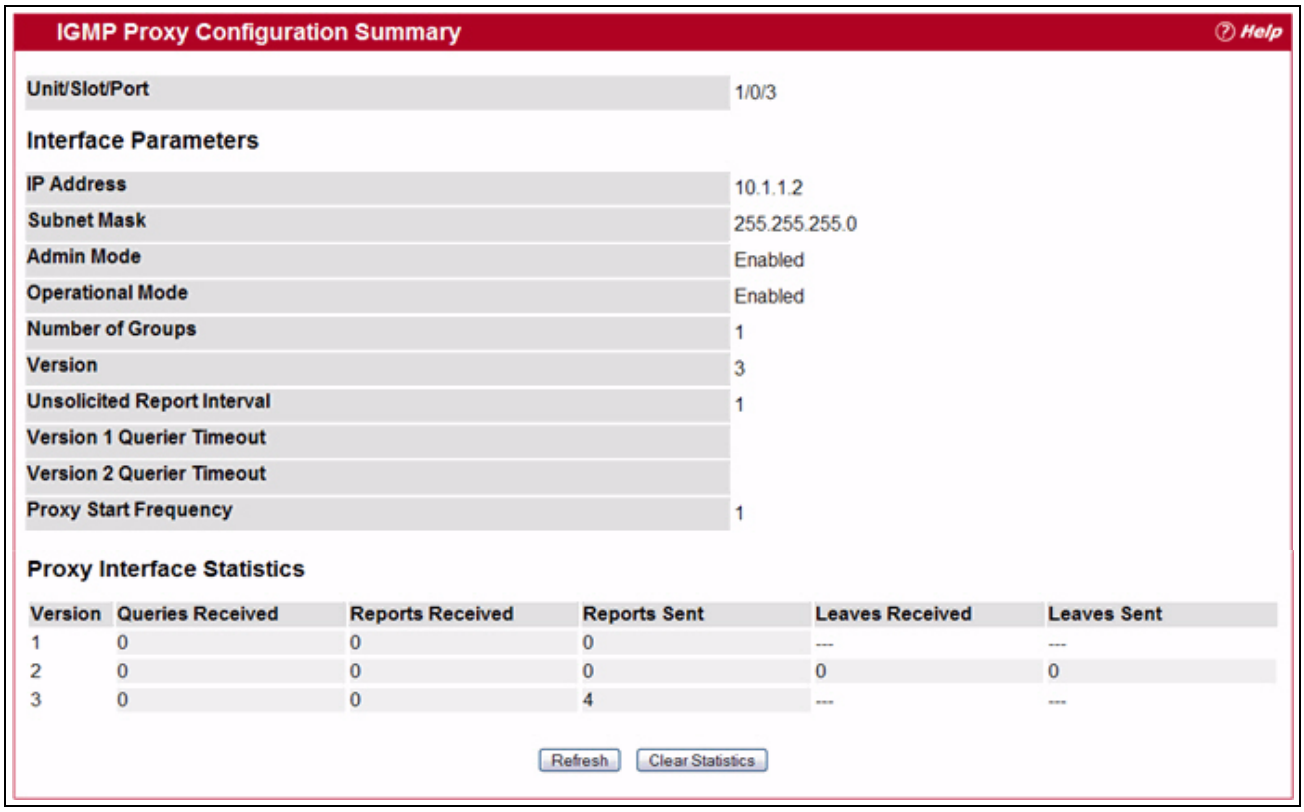

IGMP Proxy Configuration Summary

Help

Unit/Slot/Port	1/0/3

Interface Parameters

IP Address	10.1.1.2
Subnet Mask	255.255.255.0
Admin Mode	Enabled
Operational Mode	Enabled
Number of Groups	1
Version	3
Unsolicited Report Interval	1
Version 1 Querier Timeout	
Version 2 Querier Timeout	
Proxy Start Frequency	1

Proxy Interface Statistics

Version	Queries Received	Reports Received	Reports Sent	Leaves Received	Leaves Sent
1	0	0	0	---	---
2	0	0	0	0	0
3	0	0	4	---	---

Refresh Clear Statistics

Figure 7-20: IGMP Proxy Configuration Summary

Table 7-20: IGMP Proxy Configuration Summary Fields

Field	Description
Slot/Port	Displays the interface on which IGMP proxy is enabled.
IP Address	The IP address of the IGMP Proxy interface.
Subnet Mask	The subnet mask for the IP address of the IGMP Proxy interface.
Admin Mode	The administrative status of IGMP Proxy on the selected interface.
Operational Mode	The operational state of IGMP Proxy interface.
Number of Groups	The current number of multicast group entries for the IGMP Proxy interface in the cache table.
Version	The version of IGMP configured on the IGMP Proxy interface.
Unsolicited Report Interval	The Unsolicited Report Interval is the time between repetitions of a host's initial report of membership in a group. Default: 1 second. cache table.
Version 1 Querier Timeout	The older IGMP version 1 querier timeout value in seconds. The Older Version Querier Interval is the time-out for transitioning a host back to IGMPv3 mode once an older version query is heard. When an older version query is received, hosts set their Older Version Querier Present Timer to Older Version Querier Interval.
Version 2 Querier Timeout	The older IGMP version 2 querier timeout value in seconds.
Proxy Start Frequency	The number of times the proxy was brought up.
Proxy Interface Statistics	The Queries Received, Reports Received/Sent, Leaves Received/Sent are displayed in the form a table for each IGMP version.

- Click **Refresh** to refresh the data on the screen with the present state of the data in the router.
- Click **Clear Statistics** to clear the IGMP Proxy Interface statistics and reset the counters to their original values.

7.3.3.3 Interface Membership Information

This page lists each IP multicast group for which the IGMP proxy interface has received membership reports. To display this page, click **IPv4 Multicast > IGMP > Proxy Interface > Interface Membership Info** in the navigation tree.

Figure 7-21: IGMP Proxy Interface Membership Info

Table 7-21: IGMP Proxy Interface Membership Info Fields

Field	Description
Multicast Group IP	Select the IP multicast group address for which data is to be displayed. If no group membership reports have been received on the selected interface you will not be able to make this selection, and none of the non-configurable data will be displayed.
Slot/Port	Displays the interface on which IGMP proxy is enabled.
Last Reporter	The IP address of the source of the last membership report received for the IP Multicast group address on the IGMP Proxy interface.
Uptime	The time elapsed since this entry was created.
State	The state of the host entry. A Host can be in one of the state. Non-member state does not belong to the group on the interface. Delaying member state-host belongs to the group on the interface and report timer running. The report timer is used to send out the reports. Idle member state host belongs to the group on the interface and no report timer running.
Number of Sources	The number of source hosts present in the selected multicast group.

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

7.3.3.4 Interface Membership Information—Detailed

This page provides additional information on the IP multicast groups for which the IGMP proxy interface has received membership reports. To display this page, click **IPv4 Multicast > IGMP > Proxy Interface > Interface Membership Info Detailed** in the navigation tree.

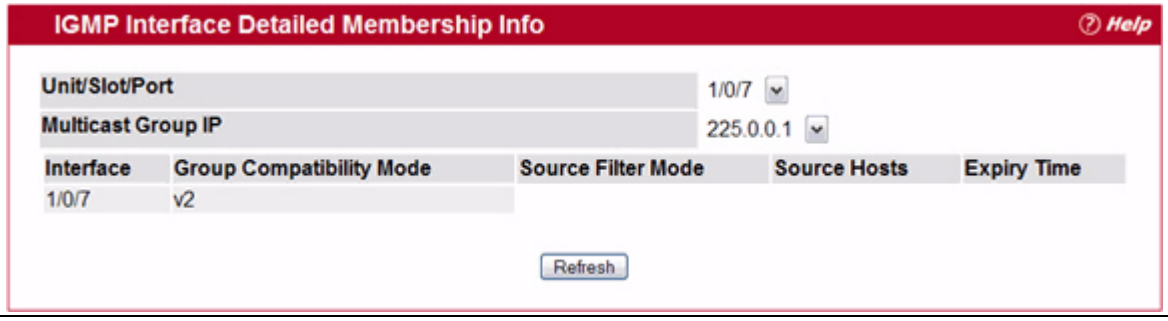

Figure 7-22: IGMP Proxy Interface Membership Info Detailed

Table 7-22: IGMP Proxy Interface Membership Info Detailed Fields

Field	Description
Multicast Group IP	Select the IP multicast group address for which data is to be displayed. If no group membership reports have been received on the IGMP Proxy interface you will not be able to make this selection, and none of the non-configurable data will be displayed.
Slot/Port	Displays the interface on which IGMP proxy is enabled.
Source IP	This parameter shows source addresses which are members of this multicast address.
Last Reporter	The IP address of the source of the last membership report received for the IP Multicast group address on the selected interface.
Up Time	Displays the up time since the entry was created in cache table.
State	The state of the host entry. A Host can be in one of the state. Non-member state - does not belong to the group on the interface. Delaying member state - host belongs to the group on the interface and report timer running. The report timer is used to send out the reports. Idle member state - host belongs to the group on the interface and no report timer running.
Filter Mode	The group filter mode (Include/Exclude/None) for the specified group on the IGMP Proxy interface.

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

7.4 Enabling and Configuring PIM-DM

Multicast protocols are used to deliver multicast packets from one source to multi-receivers. They facilitate better bandwidth utilization, and use less host and router processing, making them ideal for usage in application such as video/audio conferencing, whiteboard tools, stock distribution tickers, etc. PIM is a widely used multicast routing protocol.

Protocol Independent Multicast (PIM) is a standard multicast routing protocol that provides scalable inter-domain multicast routing across the Internet, independent of the mechanisms provided by any particular unicast routing protocol.

PIM has two types:

- PIM-Dense Mode (PIM-DM)
- PIM-Sparse Mode (PIM-SM)

PIM-DM protocol is a simple, protocol-independent multicast routing protocol. It uses existing Unicast routing table and join/prune/graft mechanism to build a tree. PIM-DM creates source-based shortest-path distribution trees making use of Reverse Path Forwarding (RPF).

PIM-DM cannot be used to build a shared distribution tree, as PIM-SM can. PIM-DM assumes that when a sender starts sending data, all downstream routers and hosts want to receive a multicast datagram. PIM-DM initially floods multicast traffic throughout the network. Routers that do not have any downstream neighbors prune back the unwanted traffic. Apart from the prune messages, PIM-DM makes use of two more messages: graft and assert. Graft messages are used whenever a new host wants to join the group. Assert messages are used to shut off duplicate flows onto the same multi-access network.

To minimize the repeated flooding of datagrams and subsequent pruning associated with a particular (S,G) pair, PIM-DM uses a State Refresh message. This message is sent by the router(s) directly connected to the source and is propagated throughout the network. When received by a router on its RPF interface, the State Refresh message causes an existing prune state to be refreshed. State Refresh messages are generated periodically by the router directly attached to the source.

There are two versions of PIM-DM. Version 2 does not use IGMP messages; instead, it uses a message that is encapsulated in IP packets with protocol number 103. In Version 2, the Hello message is introduced in place of the query message.

PIM-DM is appropriate for:

- Densely distributed receivers
- A ratio of few senders-to-many receivers (due to frequent flooding)
- High volume of multicast traffic
- Constant stream of traffic

This section describes how to configure and use PIM-DM. PIM-SM is described on page 458.

7.4.1 Global Configuration

Use this page to administratively enable or disable the PIM-DM protocol. To display the page, click **IPv4 Multicast > PIM > Global Configuration** in the navigation tree**.**

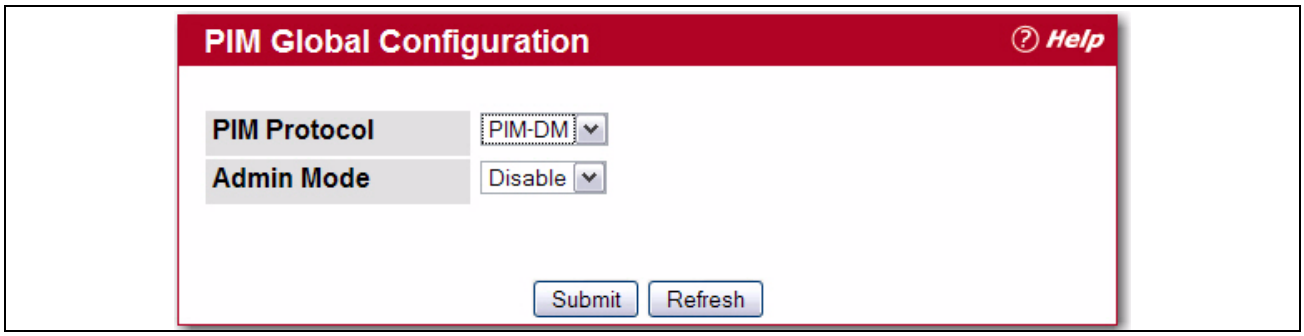

Figure 7-23: PIM Global Configuration

Table 7-23: PIM Global Configuration Fields

Field	Description
PIM Protocol Type	The protocol variant of PIM i.e Sparse mode or Dense mode to be enabled.
Admin Mode	The administrative status of PIM protocol (DM or SM as selected by PIM Protocol Type field) in the router. The default is disable.
Data Threshold Rate	The rate in K bits/second above which the last-hop router will switch to a source-specific shortest path tree. This field is applicable only for PIMSM. The valid values are from (0 to 2000) The default value is 0
Register Threshold Rate	The rate in K bits/second above which the Rendezvous Point router will switch to a source-specific shortest path tree. This field is applicable only for PIMSM. The valid values are from (0 to 2000) The default value is 0.

- Click **Submit** to send the updated configuration to the router. Configuration changes take effect immediately. These changes will not be retained across a power cycle unless a save is performed.
- Click **Refresh** to update the information on the screen.

7.4.2 Global Status

Use this page to view the PIM-DM protocol configuration. To display the page, click **IPv4 Multicast > PIM > Global Status** in the navigation tree**.**

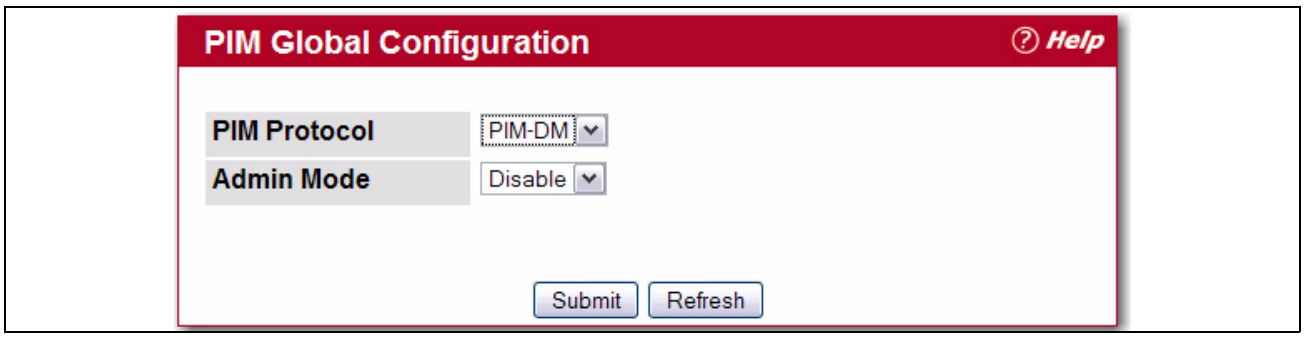

Figure 7-24: PIM Global Status

Table 7-24: PIM Global Status Fields

Field	Description
PIM Protocol Type	The protocol variant of PIM i.e Sparse mode or Dense mode to be enabled.
Admin Mode	The administrative status of PIM protocol (DM or SM as selected by PIM Protocol Type field) in the router. The default is disable.
Data Threshold Rate	The rate in K bits/second above which the last-hop router will switch to a source-specific shortest path tree. This field is applicable only for PIMSM. The valid values are from (0 to 2000) The default value is 0
Register Threshold Rate	The rate in K bits/second above which the Rendezvous Point router will switch to a source-specific shortest path tree. This field is applicable only for PIMSM. The valid values are from (0 to 2000) The default value is 0.

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

7.4.3 Interface Configuration

Use this page to configure the administrative mode of the PIM-DM protocol on an interface and to set the interval between PIM-DM messages. To display this page, click **IPv4 Multicast > PIM > Interface Configuration** in the navigation tree.

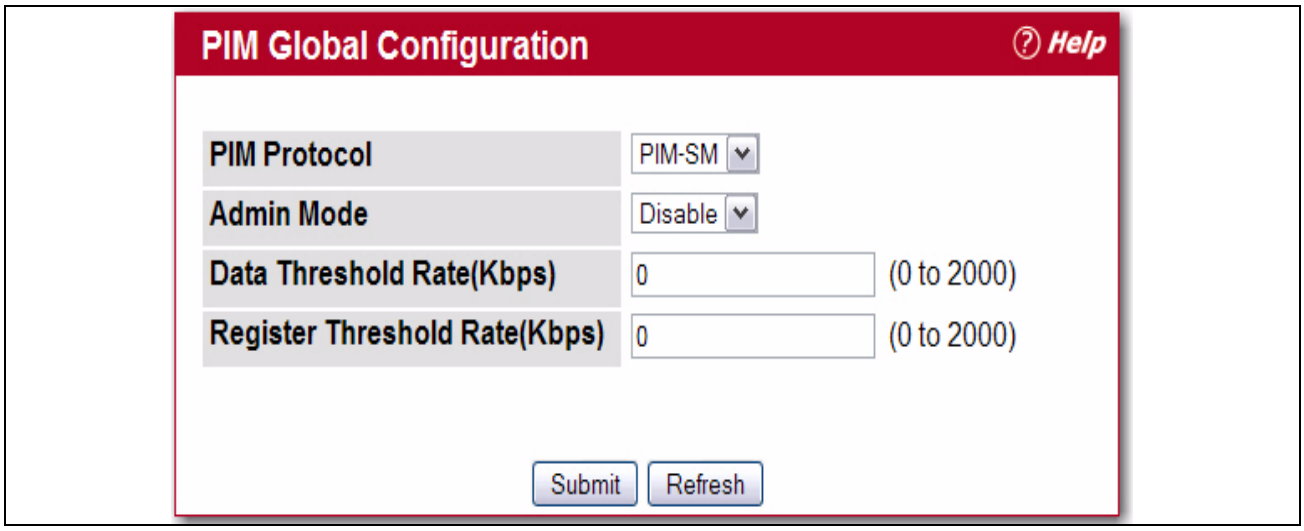

Figure 7-25: PIM-DM Interface Configuration

Table 7-25: PIM-DM Global Configuration Fields

Field	Description
Slot/Port	The slot and port for which data is to be displayed or configured. Slot 0 is the base unit.
Hello Interval	The time in seconds between the transmission of which PIM Hello messages on this interface. The valid values are from (0 to 65535). The default value is 30.
Join/Prune Interval	The frequency at which PIM Join/Prune messages are transmitted on this PIM interface. This field is applicable for PIMSM only.The valid values are from (0 to 18000). The default value is 60.
BSR Border	Sets the BSR border status on the selected interface. This field is applicable for PIMSM only.
DR Priority	Sets the DR priority for the selected interface.This field is applicable for PIMSM only. The valid values are from (0 to 2147483647) The default value is 1.

- Click **Submit** to send the updated configuration to the router. Configuration changes take effect immediately. These changes will not be retained across a power cycle unless a save is performed.
- Click **Refresh** to update the information on the screen.

7.4.4 Interface Summary

This page displays summary configuration information and operational data on the PIM-DM interface and its neighbors. To display this page, click **IPv4 Multicast > PIM > Interface Summary** in the navigation tree.

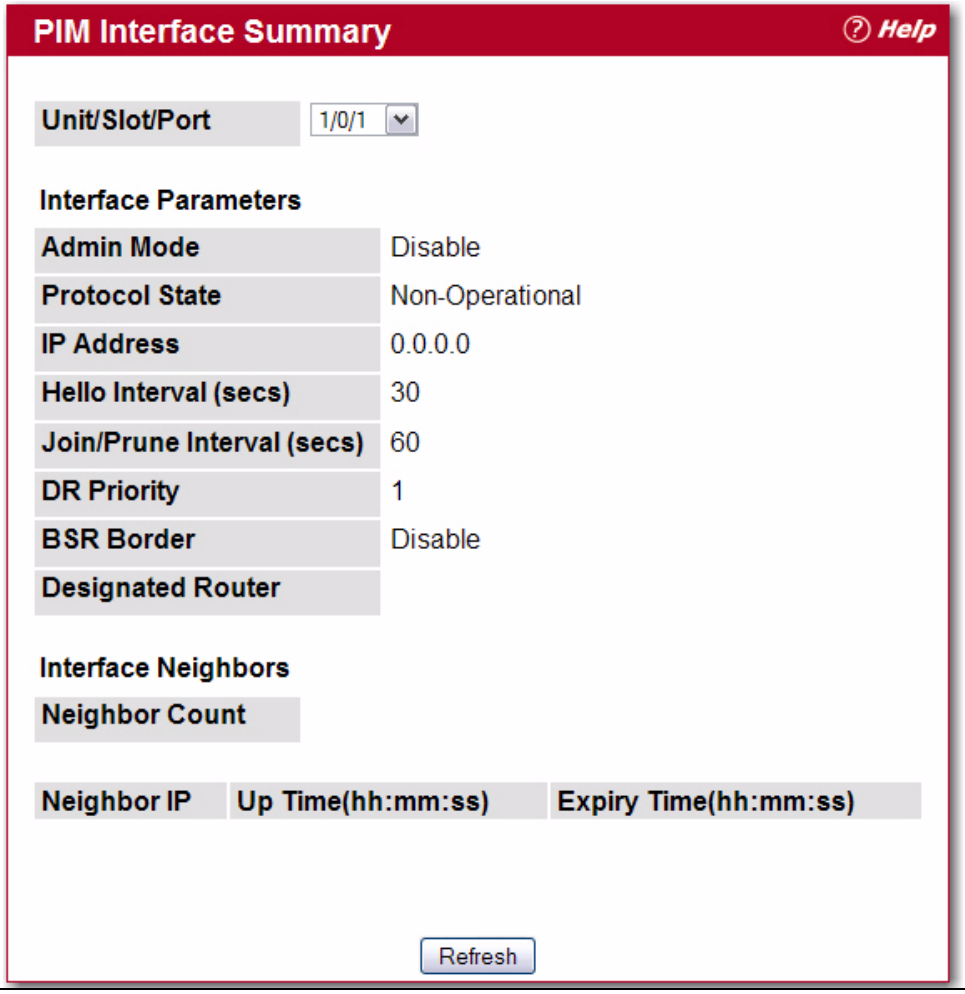

Figure 7-26: PIM Interface Summary

Table 7-26: PIM Interface Summary Fields

Field	Description
Slot/Port	The slot and port for which data is to be displayed or configured. Slot 0 is the base unit.
Admin Mode	The administrative mode of PIM-SM interface in the router (enabled or disabled).
Protocol State	The state of PIM-SM in the router (operational or non-operational).
IP Address	The IP address of the selected PIM interface.
Hello Interval	The time in seconds between the transmission of which PIM Hello messages on this interface. The valid values are from (0 to 65535). The default value is 30.
Join/Prune Interval	The frequency at which PIM Join/Prune messages are transmitted on this PIM interface. This field is applicable for PIMSM only.The valid values are from (0 to 18000). The default value is 60.
DR Priority	Sets the DR priority for the selected interface.This field is applicable for PIMSM only. The valid values are from (0 to 2147483647) The default value is 1.
BSR Border	Sets the BSR border status on the selected interface. This field is applicable for PIMSM only.
Designated Router	The designated router on the selected PIM interface. For point- to-point interfaces, this will be 0.0.0.0.

Table 7-26: PIM Interface Summary Fields (Continued)

Field	Description
Neighbor Count	The number of PIM neighbors on the selected interface.
IP Address	The IP address of the PIM neighbor for this entry.
Uptime	The time since this PIM neighbor (last) became a neighbor of the local router.
Expiry Time	The minimum time remaining before this PIM neighbor will be aged out.

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

7.5 Enabling and Configuring PIM-SM

Protocol-Independent Multicast Sparse Mode (PIM-SM) is used to efficiently route multicast traffic to multicast groups that may span wide area networks where bandwidth is a constraint.

PIM-SM uses shared trees by default and implements source-based trees for efficiency; it assumes that no hosts want the multicast traffic unless they specifically ask for it. It creates a shared distribution tree centered on a defined "rendezvous point" (RP) from which source traffic is relayed to the receivers. Senders first send the multicast data to the RP, which in turn sends the data down the shared tree to the receivers. Shared trees centered on an RP do not necessarily provide the shortest, most optimal path. In such cases PIM-SM provides a means to switch to more efficient source-specific trees. A data threshold rate is defined for toggling between trees.

PIM-SM uses a Bootstrap Router (BSR), which advertises information to other multicast routers about the rendezvous point (RP). In a given network, a set of routers can be administratively enabled as candidate bootstrap routers. If it is not apparent which router should be the BSR, the candidates flood the domain with advertisements. The router with the highest priority is elected. If all the priorities are equal, then the candidate with the highest IP address becomes the BSR.

PIM-SM is defined in RFC 4601.

7.5.1 PIM-SM Global Configuration

Use this page to configure the administrative mode of the protocol and to set threshold data rates.

To access the page, click **IPv4 Multicast > PIM > Global *Configuration*** and select **PIM-SM** from the **PIM Protocol** pulldown menu.

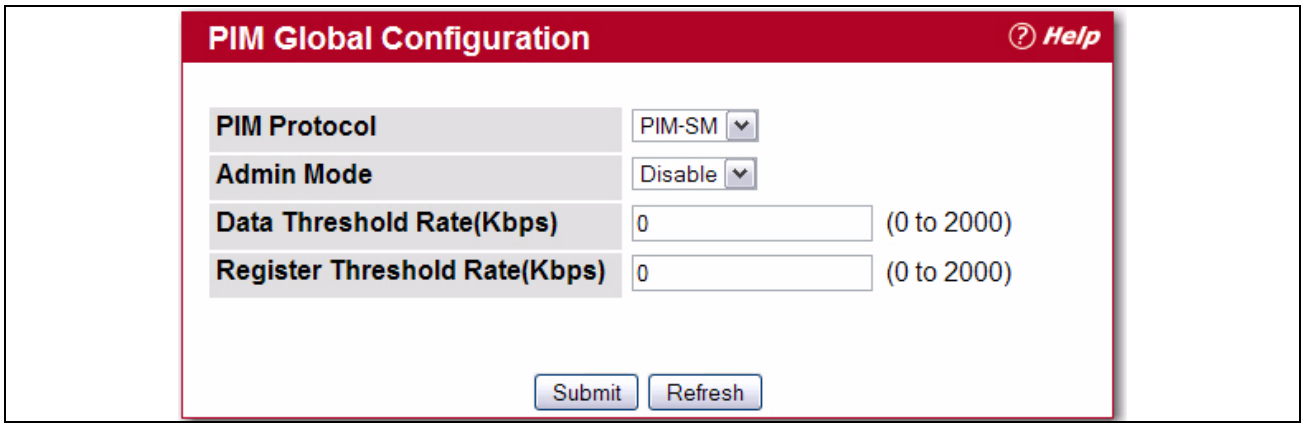

Figure 7-27: PIM-SM Global *Configuration*

Table 7-27: PIM-SM Global Configuration Fields

Field	Description
PIM Protocol Type	The protocol variant of PIM to be enabled (Sparse mode or Dense mode).
Admin Mode	Select enable or disable from the pulldown menu to set the administrative status of PIM protocol (DM or SM as selected by PIM Protocol Type field) in the router. The default is disable.
Data Threshold Rate	The rate in K bits/second above which the last-hop router will switch to a source-specific shortest path tree. This field is applicable only for PIMSM. The valid values are from (0 to 2000) The default value is 0.
Register Threshold Rate	The rate in K bits/second above which the Rendezvous Point router will switch to a source-specific shortest path tree. This field is applicable only for PIMSM. The valid values are from (0 to 2000) The default value is 0.

- Click **Submit** to send the updated configuration to the router. Configuration changes take effect immediately. These changes will not be retained across a power cycle unless a save is performed.
- Click **Refresh** to update the information on the screen.

The settings you configure on this page are also reflected on the **Global Status** page.

7.5.2 PIM-SM Global Status

Use this page to configure the administrative mode of the protocol and to set threshold data rates.

To access the page, click **IPv4 Multicast > PIM > Global Configuration** .

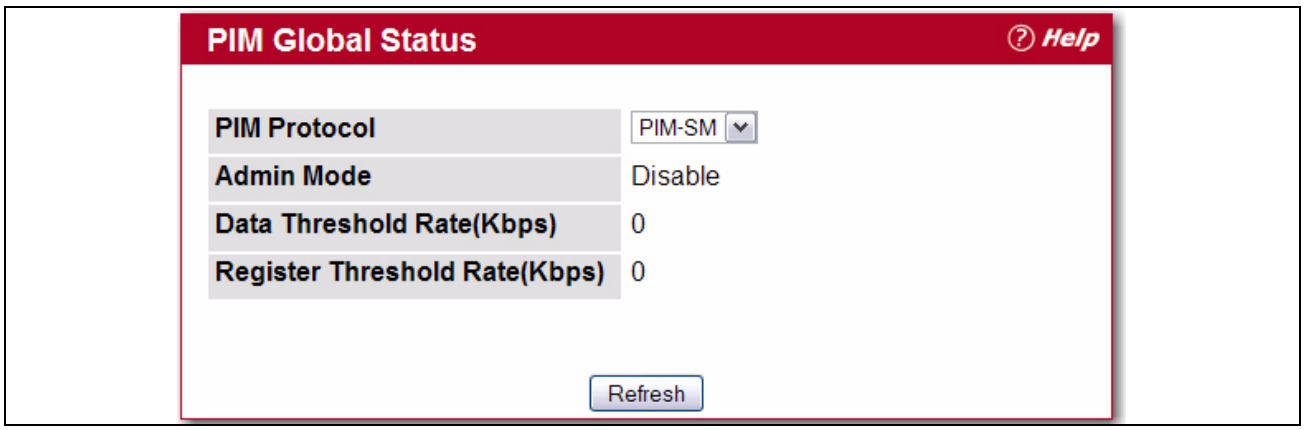

Figure 7-28: PIM-SM Global Status

Table 7-28: PIM Global Status Fields

Field	Description
PIM Protocol Type	The protocol variant of PIM to be enabled (Sparse mode or Dense mode).
Admin Mode	Select enable or disable from the pulldown menu to set the administrative status of PIM protocol (DM or SM as selected by PIM Protocol Type field) in the router. The default is disable.
Data Threshold Rate	The rate in K bits/second above which the last-hop router will switch to a source-specific shortest path tree. This field is applicable only for PIMSM. The valid values are from (0 to 2000) The default value is 0.
Register Threshold Rate	The rate in K bits/second above which the Rendezvous Point router will switch to a source-specific shortest path tree. This field is applicable only for PIMSM. The valid values are from (0 to 2000) The default value is 0.

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

7.5.3 Interface Configuration

This page configures selected interfaces to use the PIM-SM protocol. To access this page, click **IPv4 Multicast > PIM > Interface Configuration** in the navigation tree .

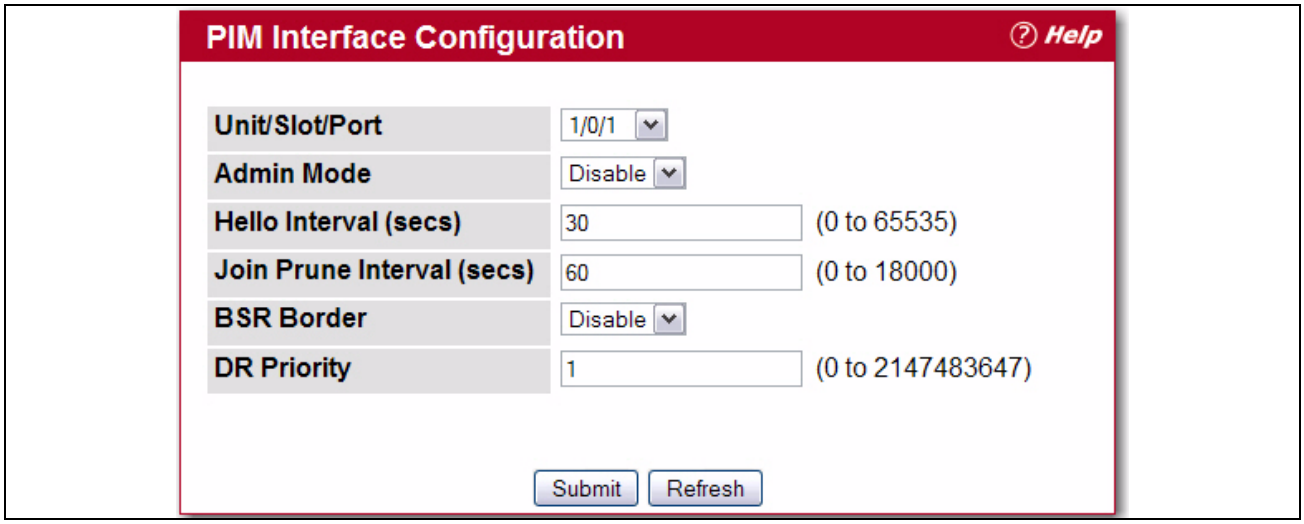

Figure 7-29: PIM Interface Configuration

Table 7-29: PIM Interface Configuration

Field	Description
Slot/Port	Select the slot and port for which data is to be displayed. Slot 0 is the base unit.
Admin Mode	Select the administrative status of PIM-SM in the router either enable or disable.

Table 7-29: PIM Interface Configuration (Continued)

Field	Description
Hello Interval	Enter the interval in seconds after which PIM Hello messages are retransmitted on the selected interface.
Join/Prune Interval	Enter the default interval in seconds at which periodic PIM-SM Join/Prune messages are to be sent.
BSR Border	Indicates whether or not this interface acts as a border for all PIM bootstrap messages. Bootstrap messages do not cross the BSR border.
DR Priority	The Designated Router priority value. The router with the highest priority value is elected as the Designated Router. A shared-media such as Ethernet may have multiple PIM-SM routers connected to it. A single one of these routers, the DR, acts on behalf of directly connected hosts with respect to the PIM-SM protocol.

After entering all required data, click **Submit** to configure an interface to use PIM-SM.

7.5.4 Interface Summary

Use this page to view PIM-SM information on a selected interface. To view the page, click **IPv4 Multicast > PIM > Interface Summary** in the navigation tree.

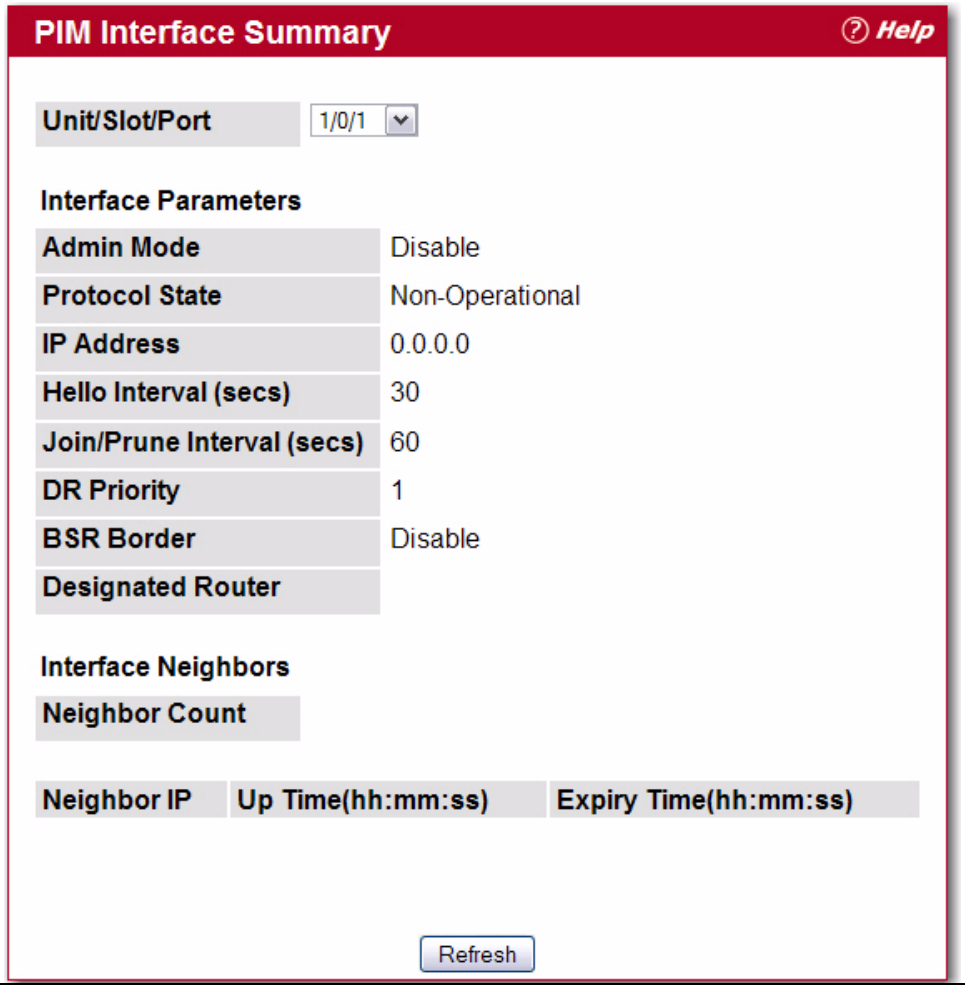

Figure 7-30: PIM Interface Configuration

Table 7-30: PIM Interface Configuration

Field	Description
Slot/Port	Select the slot and port for which data is to be displayed. Slot 0 is the base unit.
Interface Parameters	
Admin Mode	The administrative status of PIM-SM in the router either enable or disable.
Protocol State	Indicates the operational state of the PIM-SM protocol on the interface.
IP Address	The IP address of the selected PIM interface.
Hello Interval	The frequency at which PIM Hello messages are transmitted on the selected interface.
Join/Prune Interval	The frequency at which PIM Join/Prune messages are transmitted on this PIM interface. This field is always shown with defaults for PIMDM as configurability is not supported in PIMDM.
DR Priority	The DR priority on the PIM interface. This field is displayed for PIMSM only.
BSR Border	The BSR border mode on the PIM interface. This field is displayed for PIMSM only.
Designated Router	The Designated Router on the selected PIM interface. This field is displayed for PIMSM only.
Neighbor Count	The number of PIM neighbors on the selected interface.
Interface Parameters	
IP Address	The IP address of the PIM neighbor for this entry.
Up Time	The time since this PIM neighbor (last) became a neighbor of the local router.
Expiry Time	The minimum time remaining before this PIM neighbor will be aged out.

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

7.5.5 SSM Configuration

While PIM-SM employs a specially-configured RP router that serves as a meeting junction for multicast senders and listeners, Protocol-Independent Multicast Source Specific Multicast (PIM-SSM) does not use an RP. It supports only source-route deliver trees. It is used between routers so that they can track which multicast packets to forward to each other and to their directly-connected LANs.

The SSM service model can be implemented with a strict subset of the PIM-SM protocol mechanisms. Both regular IP Multicast and SSM semantics can coexist on a single router and both can be implemented using the PIM-SM protocol. A range of multicast addresses, currently 232.0.0.0/8 in IPv4, is reserved for SSM.

Use the SSM Range Configuration page to configure the SSM group IP addresses.

To access the page, click **IPv4 Multicast > PIM > SSM Configuration** in the navigation tree.

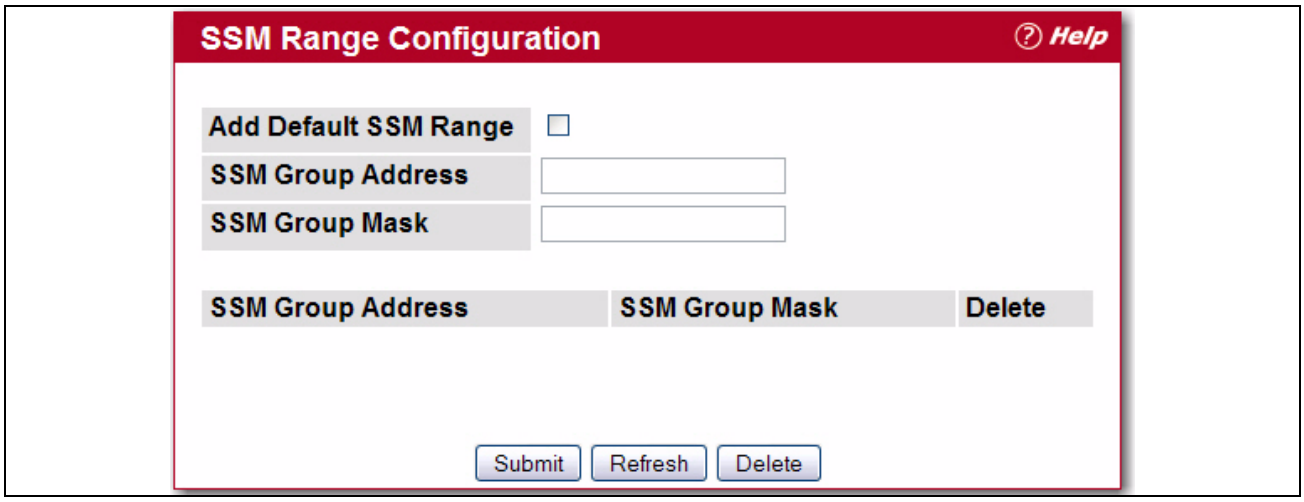

Figure 7-31: SSM Range Configuration

Table 7-31: SSM Range Configuration Fields

Field	Description
Add Default SSM Range	Select this check-box to add default SSM range (232.0.0.0/8).
SSM Group Address	Enter the IP address that, together with the mask, defines the range of available multicast addresses.
SSM Group Mask	For IPv4 configuration, enter the network mask.

- To create a new SSM group, enter the SSM group IP address and network mask and click **Submit**. The SSM groups display in the table the bottom of the page.
- Click **Refresh** to redisplay the SSM Group information with the latest data from the router.
- Click **Delete** to delete an existing SSM Group address.

7.5.6 Static RP Configuration

Use this page to configure a router as the RP for the specified group. When the RP is defined statically, it will not be chosen by election. To display this page, click **IPv4 Multicast > PIM > Static RP Configuration** in the navigation tree.

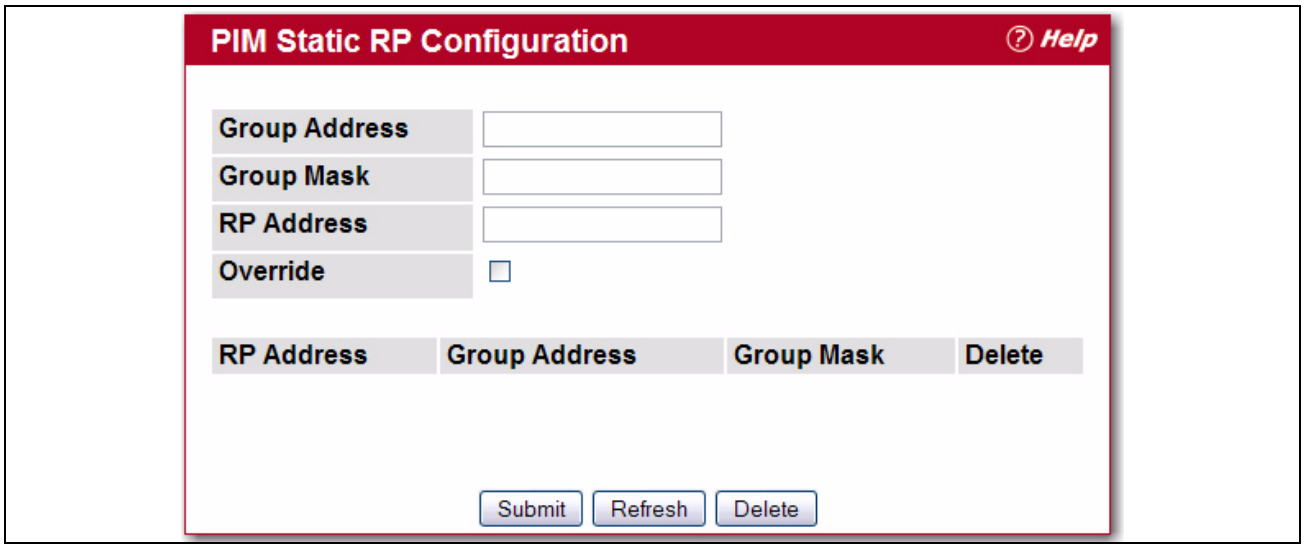

Figure 7-32: Static RP Configuration

Table 7-32: Static RP Configuration Summary

Field	Description
Group Address	Group Address for which the static RP is to be created or deleted.
Group Mask	Group Mask or which the static RP is to be created or deleted.
RP Address	IP Address of the RP for the group range created or deleted.
Override	Check to configure the static RP to override the dynamic (candidate) RPs learned for same group ranges.

- Click **Submit** to create the specified static RP Address for the PIM-SM router.
- Click **Refresh** to update the information on the screen.
- Click **Delete** to delete the static RP Address configuration.

7.5.7 Candidate RP Configuration

Use this page to configure the candidate rendezvous point (RP) for each port using PIM-SM.

To access the page, click **IPv4 Multicast > PIM > Candidate RP Configuration** in the navigation tree.

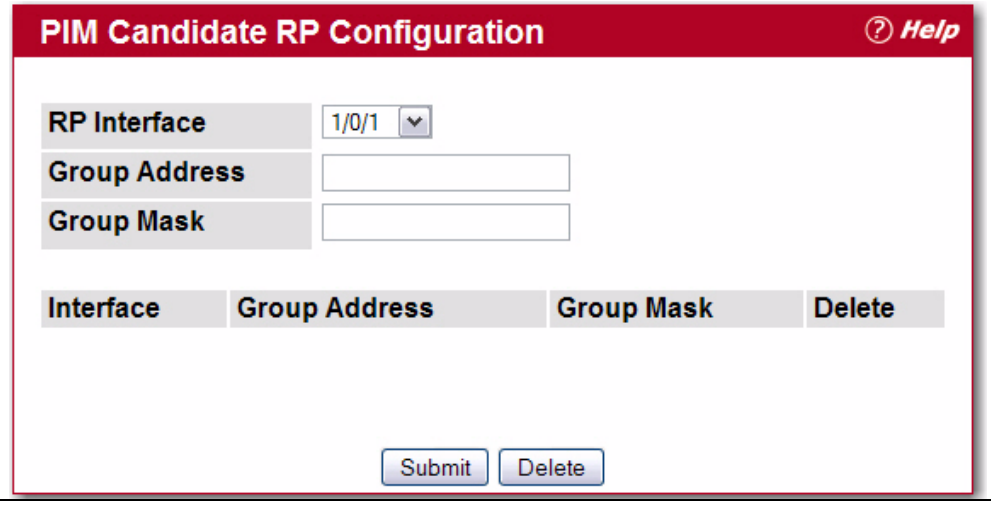

Figure 7-33: Candidate RP Configuration

Table 7-33: Candidate RP Configuration Fields

Field	Description
RP Interface	Select the slot and port for which data is to be displayed. Slot 0 is the base unit.
Group Address	The group address transmitted in Candidate-RP-Advertisements.
Group Mask (IPv4 only)	The group address mask transmitted in Candidate-RP-Advertisements.

- After entering all required data, click **Submit** to configure an interface as a PIM-SM candidate.
- To delete an RP candidate configuration, select **Delete** and click **Submit**.

7.5.8 BSR Candidate Configuration

Use this page to configure information to be used if the interface is selected as a bootstrap router. To display the page, click **IPv4 Multicast > PIM > BSR Candidate Configuration** in the navigation tree.

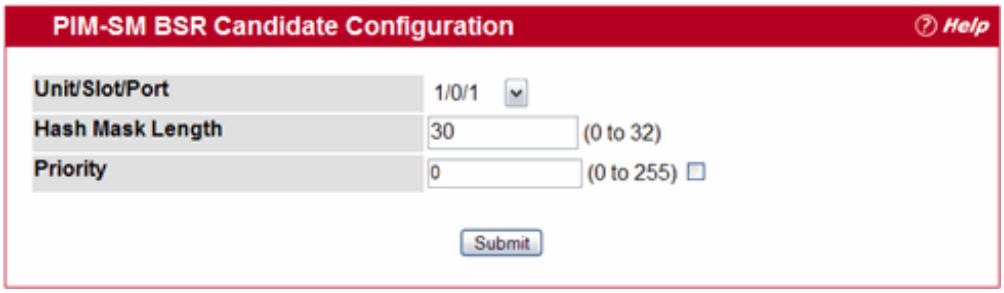

Figure 7-34: PIM-SM BSR Candidate Configuration

Table 7-34: PIM-SM BSR Candidate Configuration

Field	Description
Slot/Port	Select the slot and port for which data is to be displayed. Slot 0 is the base unit.
Hash Mask Length	The CBSR hash mask length to be advertised in bootstrap messages if this interface is elected as the bootstrap router. This hash mask length will be used in the hash algorithm for selecting the RP for a particular group.
Priority	The priority value for the local interface as a candidate bootstrap router. A value of -1 is used to indicate that the local interface is not a candidate BSR interface.

Click **Submit** to configure the interface PIM-SM BSR settings for the selected interface.

7.5.9 BSR Candidate Summary

Use this page to display information about the configured BSR candidates. To display this page, click **IPv4 Multicast > PIM > BSR Candidate Summary** in the navigation tree.

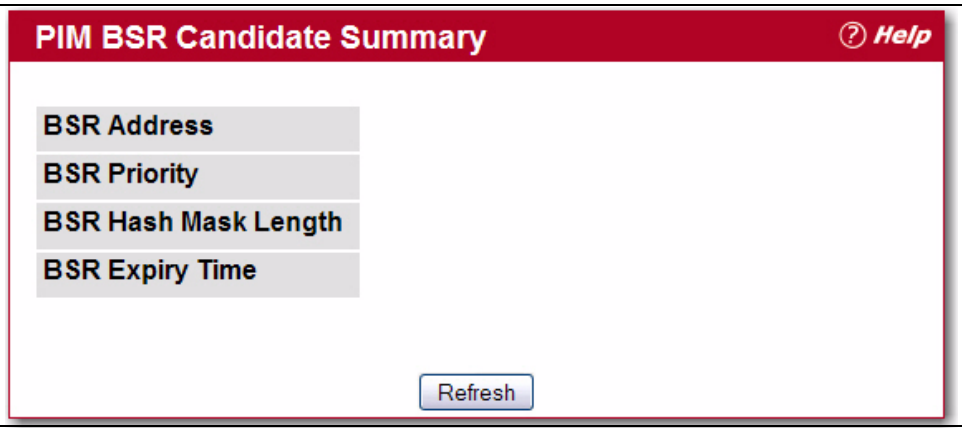

Figure 7-35: BSR Candidate Summary

Table 7-35: BSR candidate Summary

Field	Description
BSR Address	Displays the IP address of the elected bootstrap router (BSR).
BSR Priority	Displays the priority value of the elected BSR.
BSR Hash Mask Length	Displays the mask length of the elected BSR.
BSR Expiry Time	Displays the time (in hours, minutes, and seconds) in which the learnt elected bootstrap router (BSR) expires.

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

A Configuration Examples

This appendix contains examples of how to configure selected features available in the FASTPATH 6.2 ENT software. Each example contains procedures on how to configure the feature by using the Web interface, CLI, and SNMP.

This appendix describes how to perform the following procedures:

- Configuring VLANs
- Configuring VLAN Routing
- Configuring Multiple Spanning Tree Protocol
- Configuring OSPF
- Configuring 802.1x Network Access Control
- Configuring Differentiated Services for VoIP
- Configuring PIM

Note...

Each configuration example starts from a factory-default configuration unless otherwise noted.

A.1 Configuring VLANs

The diagram in this section shows a switch with four ports configured to handle the traffic for two VLANs. Port 1/0/2 handles traffic for both VLANs, while port 1/0/1 is a member of VLAN 2 only, and ports 1/0/3 and 1/0/4 are members of VLAN 3 only.

The following examples show how to create VLANs, assign ports to the VLANs, and assign a VLAN as the default VLAN to a port.

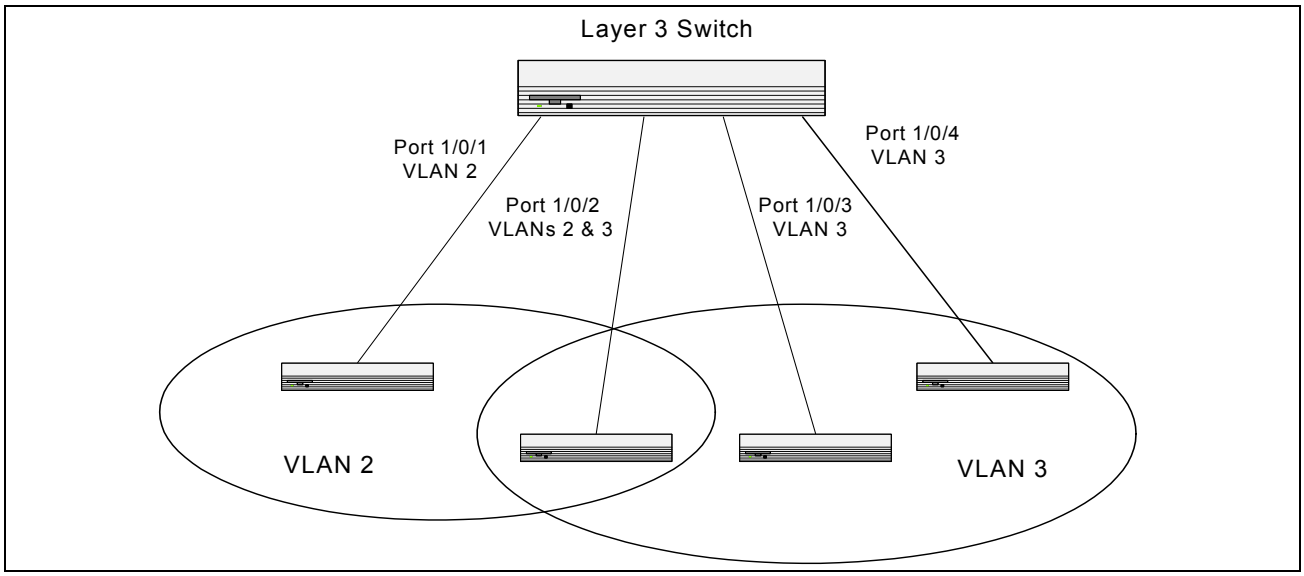

Figure A-1: VLAN Example Network Diagram

www.kontron.com

A.1.1 Using the Web Interface to Configure VLANs

1. Access the **Switching > VLAN > Configuration** page.
2. Select the Create option in the VLAN field.
3. Select the VLAN ID-Range option and enter 2 to 3 in the range fields.

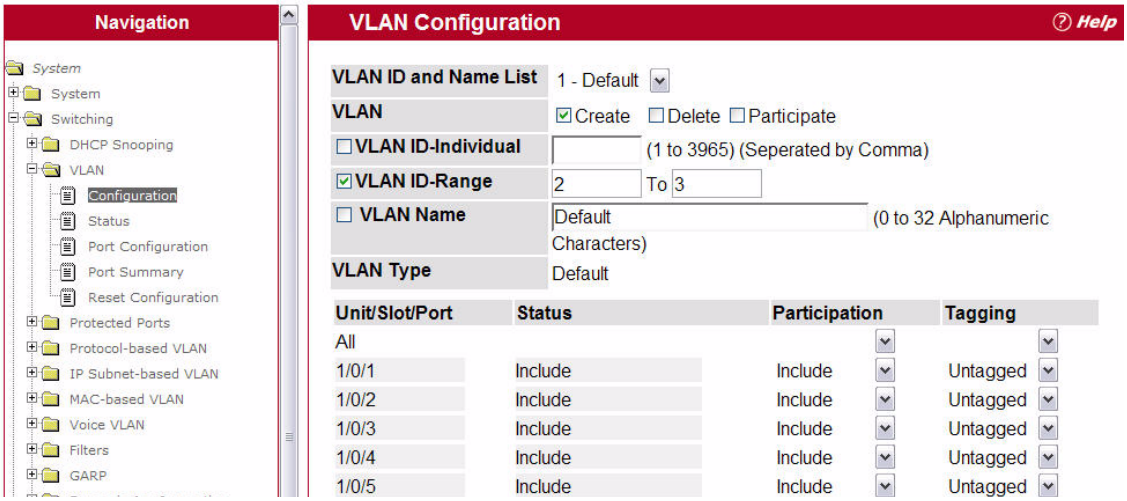

4. Click **Submit**.
5. Select VLAN 2 from the VLAN ID and Name List.
6. Select the Participate option in the VLAN field.
7. For ports 1/0/1 and 1/0/2, select Include from the Participation menu to specify that these ports are members of VLAN 2.
8. From the Tagging menu, select Tagged in the first row (All) to specify that frames will always be transmitted tagged from ports that are members of VLAN 2.

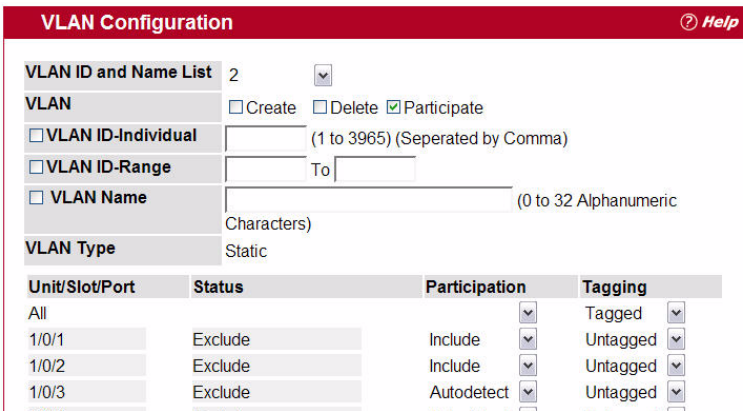

9. Click **Submit**.
10. Select VLAN 3 from the VLAN ID and Name List.
11. Select the Participate option in the VLAN field.
12. For ports 1/0/2, 1/0/3 and 1/0/4, select Include from the Participation menu to specify that these ports are members of VLAN 3.
13. Click **Submit**.
14. Go to the **Switching > VLAN > Port Configuration** page.
15. From the Slot/Port menu, select 0/1.
16. In the Acceptable Frame Types field, select VLAN Only to specify that untagged frames will be rejected on receipt.

17. Click **Submit**.

18. From the Slot/Port menu, select 0/2.

19. In the Port VLAN ID field, enter 3 to assign VLAN 3 as the default VLAN for the port.

20. In the Acceptable Frame Types field, select VLAN Only to specify that untagged frames will be rejected on receipt.

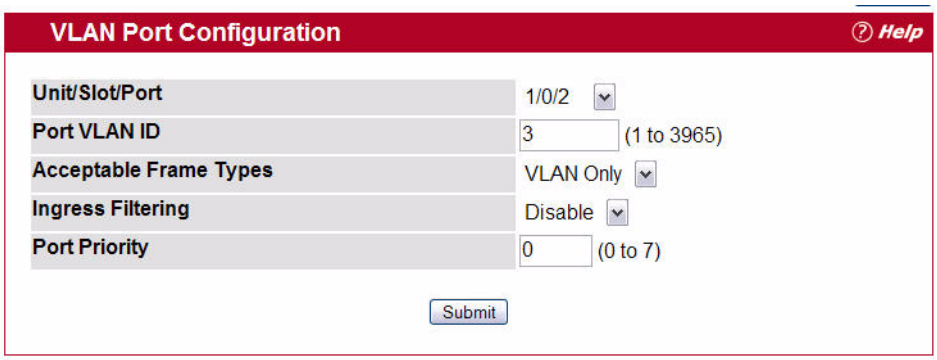

21. Click **Submit.**

A.1.2 Using the CLI to Configure VLANs

1. Create VLAN 2 and VLAN 3.

```
(Broadcom FASTPATH Routing) #vlan database
   vlan 2
   vlan 3
   exit
```

2. Assign ports 1/0/1 and 1/0/2 to VLAN2 and specify that untagged frames will be rejected on receipt.

```
(Broadcom FASTPATH Routing) #Config
   interface 1/0/1
      vlan participation include 2
      vlan acceptframe vlanonly
      exit
   interface 1/0/2
      vlan participation include 2
      vlan acceptframe vlanonly
```

3. While in interface config mode for port 1/0/2, assign VLAN3 as the default VLAN.

```
   (Broadcom FASTPATH Routing) (Interface 1/0/2)#vlan pvid 3
   exit
```

4. Specify that frames will always be transmitted tagged from ports that are members of VLAN 2.

```
Config
   vlan port tagging all 2
   exit
```

5. Assign the ports that will belong to VLAN 3.

> **Note...**
>
> Port 1/0/2 belongs to both VLANs, and port 1/0/1 can never belong to VLAN 3.

```
(Broadcom FASTPATH Routing) #Config
   interface 1/0/2
      vlan participation include 3
      exit
   interface 1/0/3
      vlan participation include 3
      exit
```

```
    interface 1/0/4
       vlan participation include 3
       exit
    exit
```

6. Specify that untagged frames will be accepted on port 1/0/4.

```
 (Broadcom FASTPATH Routing) #Config
    interface 1/0/4
       vlan acceptframe all
       exit
    exit
```

A.1.3 Using the SNMP to Configure VLANs

1. Use the objects in dot1qVlanStaticTable (in dot1qVlan in the QBRIDGE-MIB module) to create VLANs 2 and 3.

   Set the dot1qVlanStaticRowStatus object to 'CreateandGo (4)' to create a VLAN. If the other parameters are not specified, simply specifying the dot1qVlanIndex and dot1qVlanStaticRowStatus is sufficient to create the VLAN.

   The full path to the object is iso(1).org(3).dod(6).internet(1).mgmt(2).mib-2(1).dot1dBridge(17).qBridgeMIB(7).qBridgeMIBObjects(1).dot1qVlan(4).dot1qVlanStaticTable(3).dot1qVlanStaticEntry(1).dot1qVlanStaticRowStatus(5).

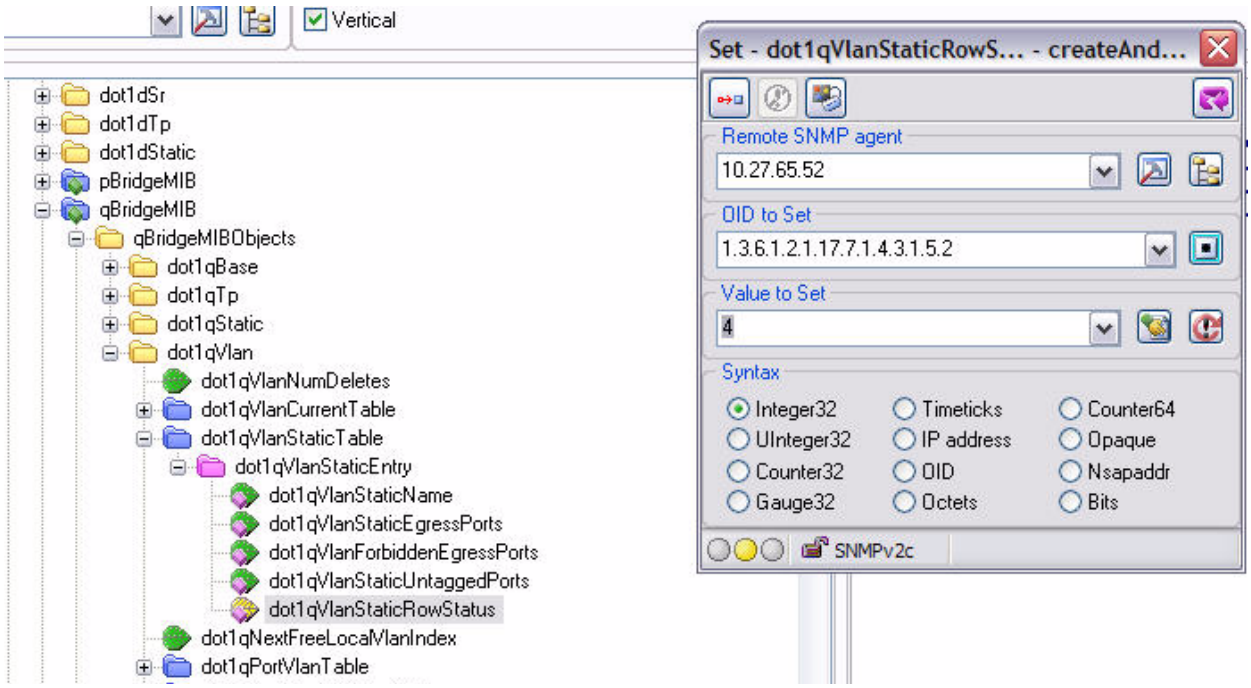

2. To assign ports 1/0/1 and 1/0/2 to VLAN2, retrieve the current dot1qStaticEgressPorts mask and append interfaces 1/0/1 and 1/0/2 to this mask by setting the first octet to 0xC0.

   The dot1qVlanStaticEgressPorts bit mask can be constructed according to the following rules:

   - Each octet within this value specifies a set of eight ports, with the first octet specifying ports (1-8), the second octet specifying ports (9-16), and so on.
   - Within each octet, the most significant bit represents the lowest numbered port, and the least significant bit represents the highest numbered port. Thus, each port of the bridge is represented by a single bit within the value of this object. If that bit has a value of (1), then that port is included in the set of ports. The port is not included if its bit has a value of (0).

   For example if the switch has 12 ports and we want to add ports 1 and 4 in the VLAN and exclude all other ports, then the bit mask in hex will be 0x50 0x00.

3. To specify that frames will always be transmitted tagged from ports that are members of VLAN 2, use the dot1qVlanStaticUntaggedPorts object and set the value of the appropriate number of octets to 0.

   Each octet represents eight ports, so for a 48-port switch, the first six octets would be zero.

4. To specify that ports 1/0/1 and 1/0/2 will only accept tagged frames and will reject untagged frames on receipt, set the dot1qPortAcceptableFrameTypes object to admitOnlyVlanTagged(2).

   The object is in dot1qPortVlanEntry in the dot1qPortVlanTable.

5. To assign VLAN3 as the default VLAN for interface 1/0/2., set the value of dot1qPvid for 1/0/2 (instance 2) to 3.
6. To assign ports 1/0/2, 1/0/3, and 1/0/4 to VLAN3, retrieve the current dot1qStaticEgressPorts mask and append the interfaces to this mask by setting the first octet to 0x70.

A.2 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.

A.2.1 Using the Web UI to Configure MSTP

1. Create VLANs 10 and 20.
   a. Access the **Switching > VLAN > Configuration** page.
   b. Select the Create option in the VLAN field.
   c. Select the VLAN ID-Individual option and enter 10.
   d. Click **Submit**.
   e. Repeat the steps to add VLAN 20.
2. Enable MSTP on the switch and change the configuration name.
3. Changing the configuration name allows all the bridges that want to be part of the same region to join.
   a. Go to the **Switching > Spanning Tree > Switch Configuration/Status** page.
   b. From the STP Mode menu, select Enable.
   c. In the Configuration Name field, enter broadcom.
   d. Click **Submit**.

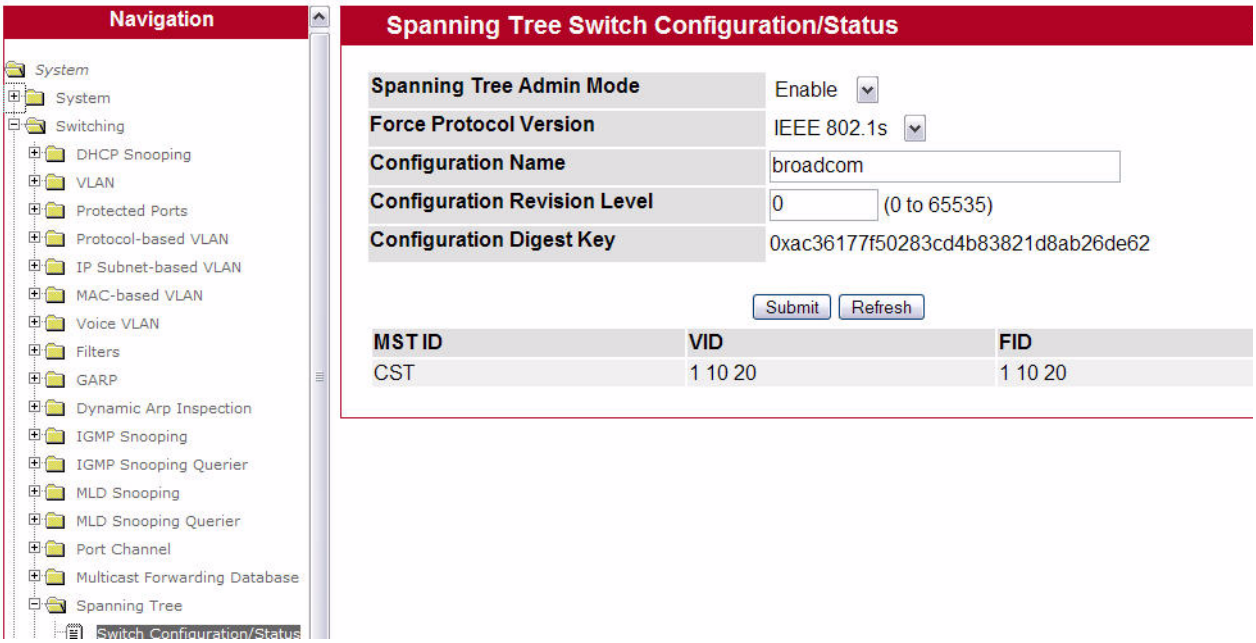

4. Create two MST instances.

a. Go to the **Switching > Spanning Tree > MST Configuration/Status** page.

b. From the MST field, select Create.

c. In the MST ID field, enter 10.

d. Click **Submit**.

e. Repeat the steps to create an MST instance with an ID of 20.

5. Associate MST ID 10 with VLAN 10 and assign a bridge priority of 16384

a. Select MST 10 from the MST menu.

b. Enter 16384 in the Bridge Priority field.

c. Click VLAN 10 to select it from the VLAN ID field.

d. Click **Submit**.

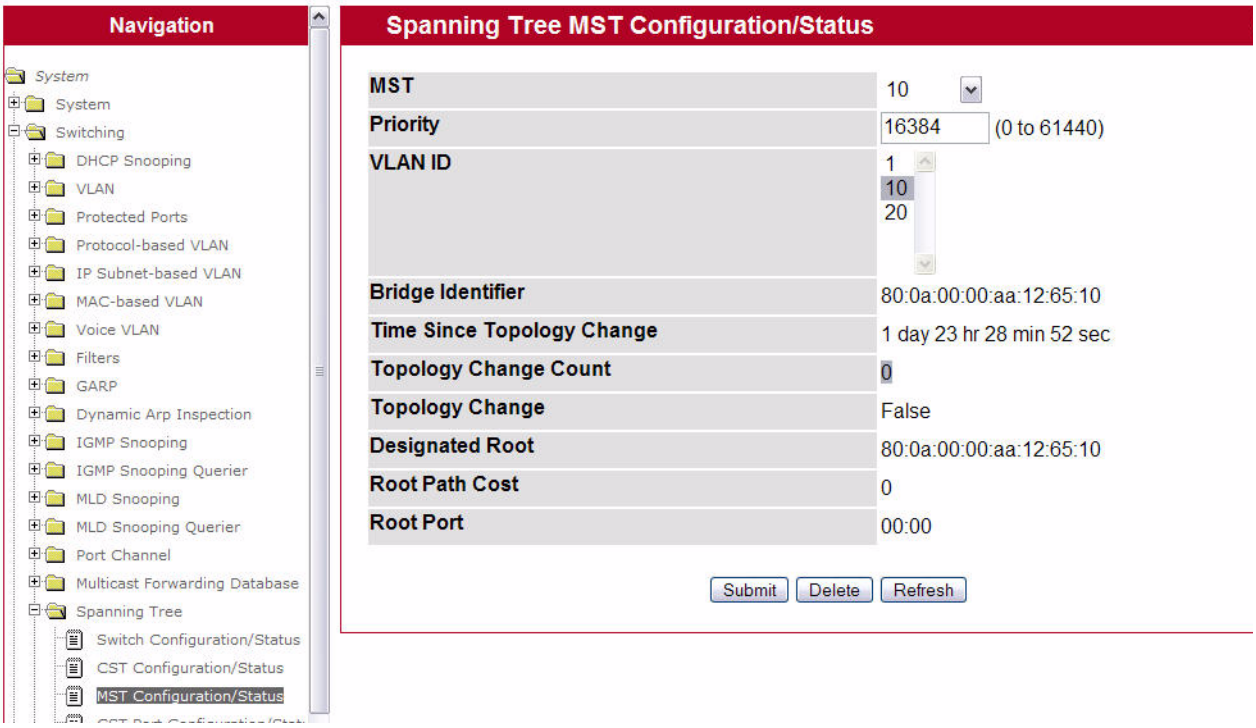

6. Use similar procedures to associate MST instance 20 to VLAN 20 and assign it a bridge priority value of 61440.

   By using a lower priority for MST 20, MST 10 becomes the root bridge.

7. Enable STP on port 1/0/1.
   a. Go to the **System > Port > Configuration** page.
   b. From the Slot/Port mode, select port 0/1.
   c. From the STP Mode menu, select Enable.
   d. Click **Submit**.

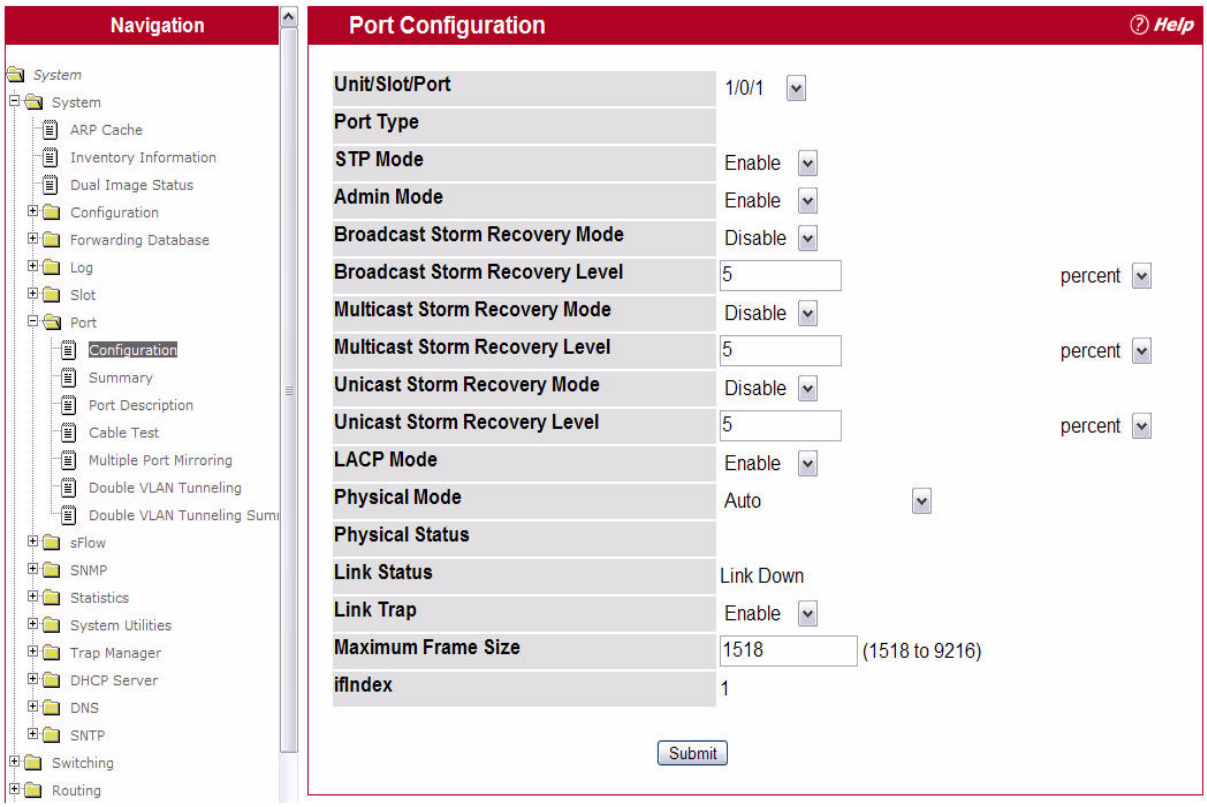

8. Use similar procedures to enable STP on port 1/0/2.
9. Force port 1/0/2 to be the root port for MST 20, which is the non-root bridge.
   a. Go to the **Switching > Spanning Tree > MST Port Configuration/Status** page.
   b. From the MST ID menu, select 20.
   c. From the Slot/Port menu, select 0/2.
   d. In the Port Priority field, enter 64.
   e. Click **Submit**.

A.2.2 Using the CLI to Configure MSTP

1. Create VLAN 10 and VLAN 20.
```
(Broadcom FASTPATH Routing) #vlan database
    vlan 10
    vlan 20
    exit
```
2. Enable spanning tree Globally
```
(Broadcom FASTPATH Routing) #config
    spanning-tree
```
3. Create MST instances 10 and 20.
```
    spanning-tree mst instance 10
    spanning-tree mst instance 20
```
4. Associate MST instance 10 to VLAN 10 and MST instance 20 to VLAN 20
```
    spanning-tree mst vlan 10 10
    spanning tree mst vlan 20 20
```
5. Change the name so that all the bridges that want to be part of the same region can form the region.
```
    spanning-tree configuration name broadcom
```

6. Make the MST ID 10 bridge the root bridge by lowering the priority.
```
spanning-tree mst priority 10 16384
```
7. Change the priority of MST ID 20 to ensure the other bridge is the root bridge.
```
spanning-tree mst priority 20 61440
```
8. Enable STP on interface 1/0/1
```
interface 1/0/1
   spanning-tree port mode
   exit
```
9. Enable STP on interface 1/0/2
```
interface 1/0/2
   spanning-tree port mode
```
10. On the non-root bridge change the priority to force port 1/0/2 to be the root port.
```
spanning-tree mst 20 port-priority 64
exit
```

A.2.3 Using SNMP to Configure MSTP

1. Use the objects in dot1qVlanStaticTable (in dot1qVlan in the QBRIDGE-MIB module) to create VLANs 10 and 20.
2. To enable spanning tree globally, set the agentStpAdminMode object in the FASTPATH-SWITCHING-MIB module to enable (2).

   The full path to the object is iso(1).org(3).dod(6).internet(1).private(4).enterprises(1).broadcom(4413).broadcomProducts(1).fastPath(1).fastPathSwitching(1).agentConfigGroup(2).agentStpSwitchConfigGroup(15).agentStpAdminMode(6).
3. Use the agentStpConfigName object in the agentStpSwitchConfigGroup to change the name so that all the bridges that want to be part of the same region can form the region.
4. Use the agentStpMstRowStatus object in the agentStpMstTable to create MST instances 10 and 20.

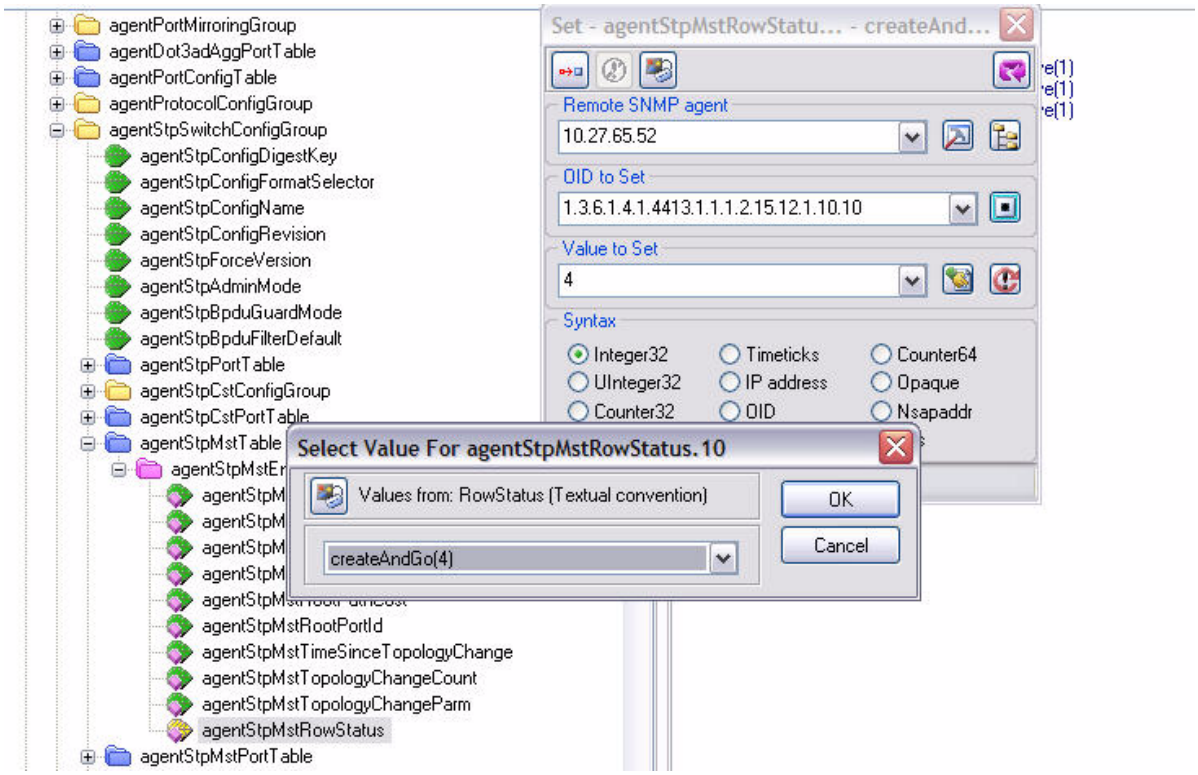

5. Use the agentStpMstBridgePriority object to set the bridge priorities for MST 10 and MST 20:
   - For MST ID 10, set the value to 16384 to make it the root bridge.
   - For MST ID 20, set the value to 61440 to ensure the other bridge is the root bridge.
6. Use the agentStpMstVlanRowStatusAssociate object in the agentStpMstVlanTable to associate MST instance 10 to VLAN 10 and MST instance 20 to VLAN 20.
   - For MST ID 20, the OID to set is 1.3.6.1.4.1.4413.1.1.1.2.15.14.1.1.10.10 (the final .10 is the VLAN ID)
   - For MST ID 20, the OID to set is 1.3.6.1.4.1.4413.1.1.1.2.15.14.1.1.20.20

   Set the value to CreateAndGo (4)
7. Use the agentStpPortState in agentStpPortTable under agentStpSwitchConfigGroup to enable STP on interface 1/0/1 and interface 1/0/2.

   For instance 1 and 2, set the value to enable (1).
8. Use the agentStpMstPortPriority object in agentStpMstPortTable to change the port priority on interface 1/0/2 to force the port to be the root port on the non-root bridge.

   For instance 2, set the value to 64.

A.3 Configuring VLAN Routing

This section provides an example of how to configure FASTPATH 6.2 ENT 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 FASTPATH 6.2 ENT software to provide the VLAN routing support shown in the diagram.

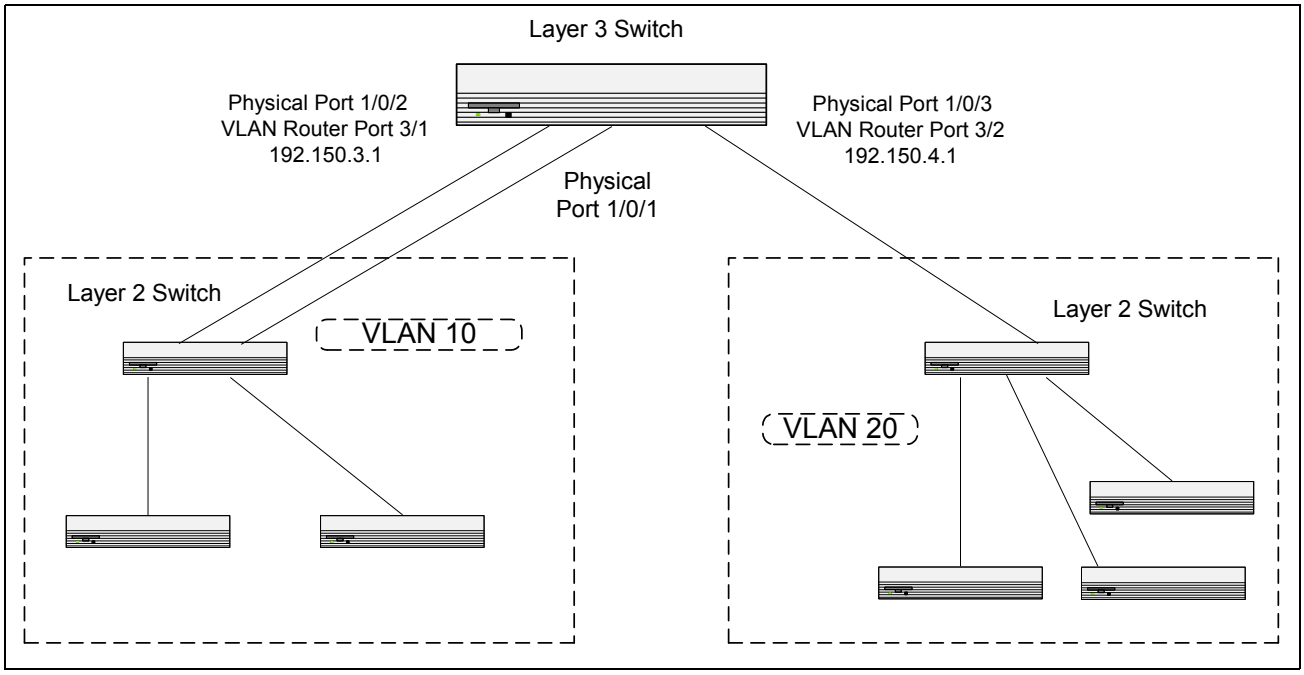

Figure A-2: VLAN Routing Example Network Diagram

A.3.1 Using the Web Interface to Configure VLAN Routing

Use the following screens to perform the same configuration using the Web Interface:

1. From the **Switching > VLAN > Configuration** page, perform the following configuration:
   - Create VLANs 10 and 20.
   - Include interfaces 1/0/1 and 1/0/2 as members of VLAN 10, and set tagging for all interfaces to Tagged.
   - Include interface 1/0/3 as a member of VLAN 20, and set tagging for all interfaces to Tagged.
2. From the **Switching > VLAN > Port Configuration** page, set the port VLAN ID for interfaces 1/0/1 and 1/0/2 to 10 and the port VLAN ID for interface 1/0/3 to 20.
3. Navigate to the **Routing > VLAN Routing > Configuration** page.
4. Enter 10 in the VLAN ID field, and then click **Create**.

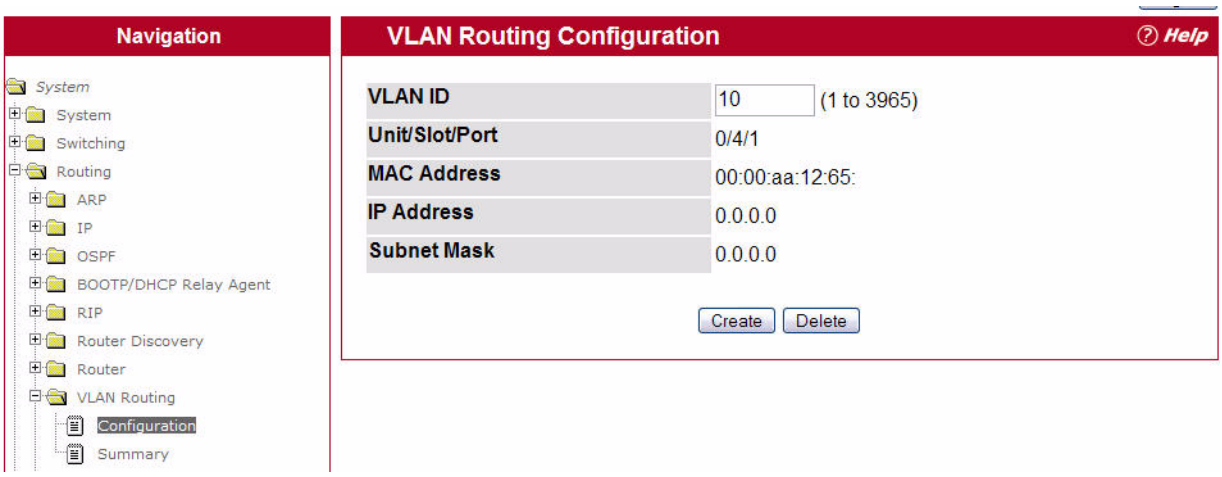

5. Enter 20 in the VLAN ID field, and then click **Create**.
6. Go to the **Routing > VLAN Routing > Summary** page to view the logical interface IDs assigned to the VLAN routing interfaces.

   VLAN 10 is assigned ID 0/4/1 and VLAN 20 is assigned ID 0/4/2
7. To enable routing on the switch, go to the **Routing > IP > Configuration** page, select Enable from the Routing Mode menu, and click **Submit**.

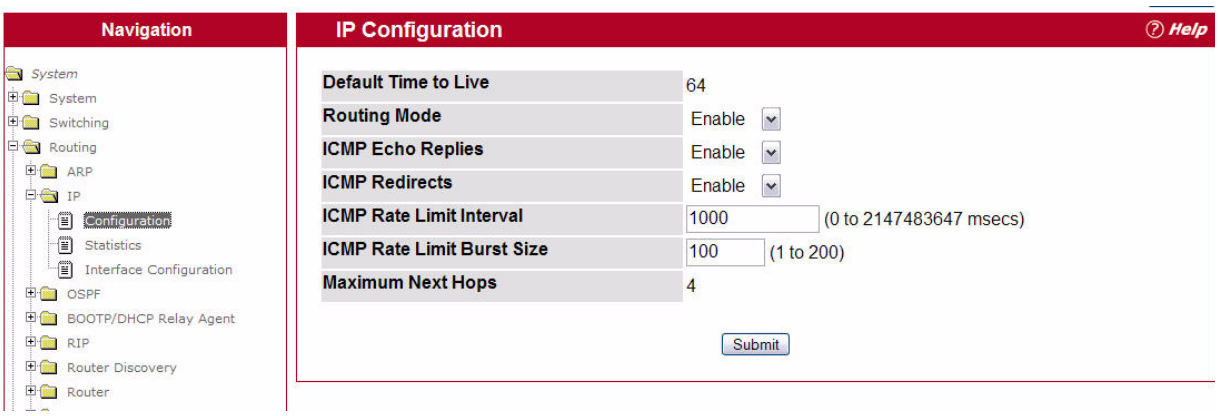

8. Go to the **Routing > IP > Interface Configuration** page to configure the IP addresses and subnet masks for the virtual router ports.
   a. From the Slot/Port menu, select 4/1.
   b. Enter 192.150.3.1 in the IP Address field.
   c. Enter 255.255.255.0 in the Subnet Mask field.
   d. Click **Submit**.

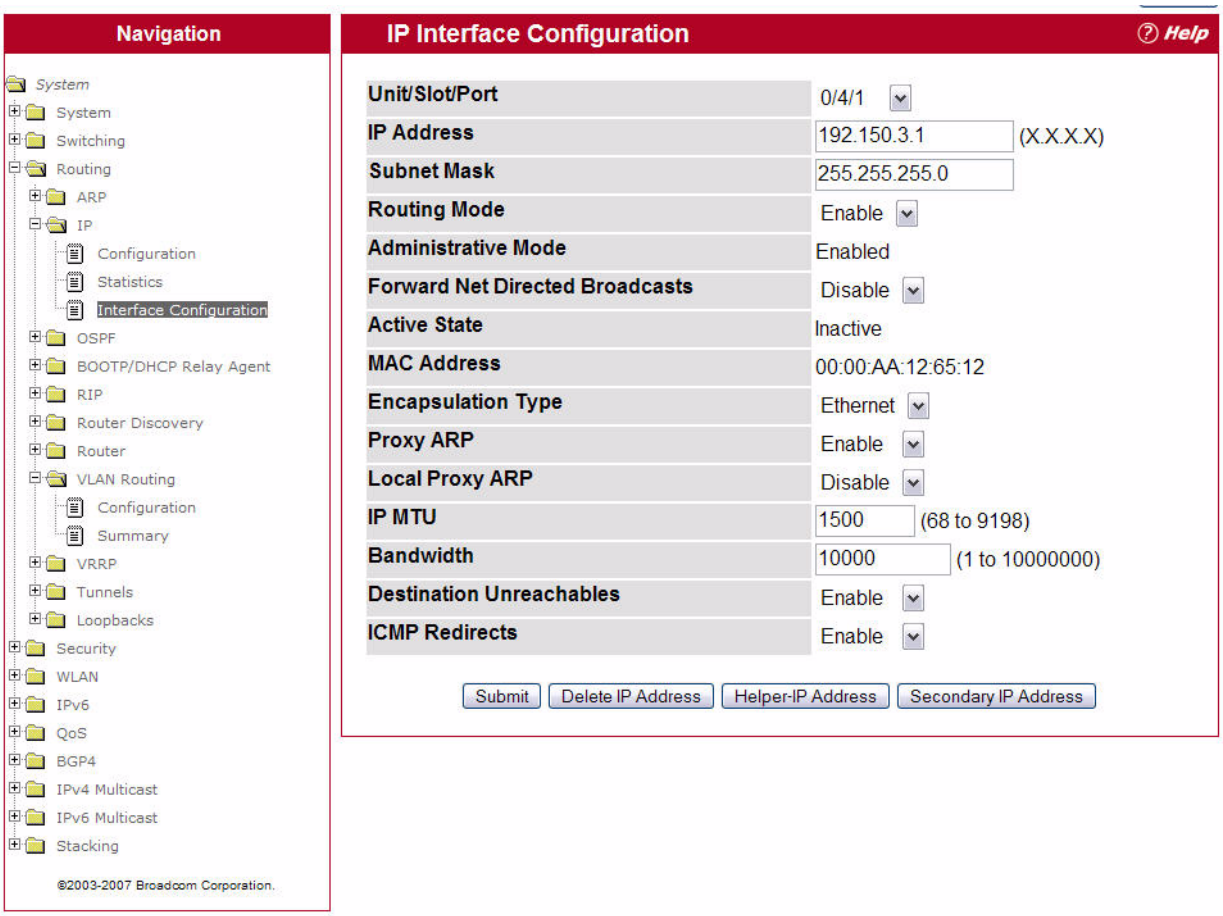

Select interface 0/4/2 from the Unit/Slot Port menu and configure it with an IP address of 192.150.4.1 and subnet mask of 255.255.255.0.

A.3.2 Using the CLI to Configure VLAN Routing

1. Create VLAN 10 and VLAN 20.

```
(Broadcom FASTPATH Routing) #vlan database
   vlan 10
   vlan 20
   exit
```

2. Configure ports 1/0/1, 1/0/2 as members of VLAN 10 and specify that untagged frames received on these ports will be assigned to VLAN 10.

```
config
   interface 1/0/1
      vlan participation include 10
      vlan pvid 10
      exit
   interface 1/0/2
      vlan participation include 10
```

```
        vlan pvid 10
        exit
```

3. Configure port 1/0/3 as a member of VLAN 20 and specify that untagged frames received on these ports will be assigned to VLAN 20

```
    interface 1/0/3
        vlan participation include 20
        vlan pvid 20
        exit
    exit
```

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:

```
 (Broadcom FASTPATH Routing) #vlan database
    vlan routing 10
    vlan routing 20
    exit
```

6. View the logical interface IDs assigned to the VLAN routing interfaces.

```
 (Broadcom FASTPATH Routing) #show ip vlan

 MAC Address used by Routing VLANs:   00:00:AA:12:65:12

           Logical
 VLAN ID   Interface        IP Address       Subnet Mask
 -------  --------------  ---------------  ---------------
 10       0/4/1           0.0.0.0          0.0.0.0
 20       0/4/2           0.0.0.0          0.0.0.0
```

As the output shows, VLAN 10 is assigned ID 0/4/1 and VLAN 20 is assigned ID 0/4/2

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 0/4/1
        ip address 192.150.3.1 255.255.255.0
    exit
    interface 0/4/2
        ip address 192.150.4.1 255.255.255.0
        exit
    exit
```

A.3.3 Using SNMP to Configure VLAN Routing

1. Use the dot1qVlanStaticRowStatus object in the dot1qVlanStaticTable to create VLAN 10 and VLAN 20.
2. To configure VLAN membership, retrieve the current dot1qStaticEgressPorts mask and append the desired interfaces to the mask.
   - VLAN 10: 1/0/1 and 1/0/2
   - VLAN 20: 1/0/3
3. To assign the PVID for an interface, use the dot1qPvid object.
   - 1/0/1: PVID 10
   - 1/0/2: PVID 10
   - 1/0/3: PVID 20

4. To specify that all frames transmitted for VLANs 10 and 20 will be tagged, use the dot1qVlanStaticUntaggedPorts object and set the value of the appropriate number of octets to 0.

   Each octet represents eight ports, so for a 48-port switch, the first six octets would be zero.

5. To enable routing for the VLANs, use the agentSwitchIpVlanRoutingStatus object in the agentSwitchIpVlanTable under agentSwitchIpGroup in fastPathRouting to set the value for VLAN 10 and VLAN 20 to CreateAndGo (4).

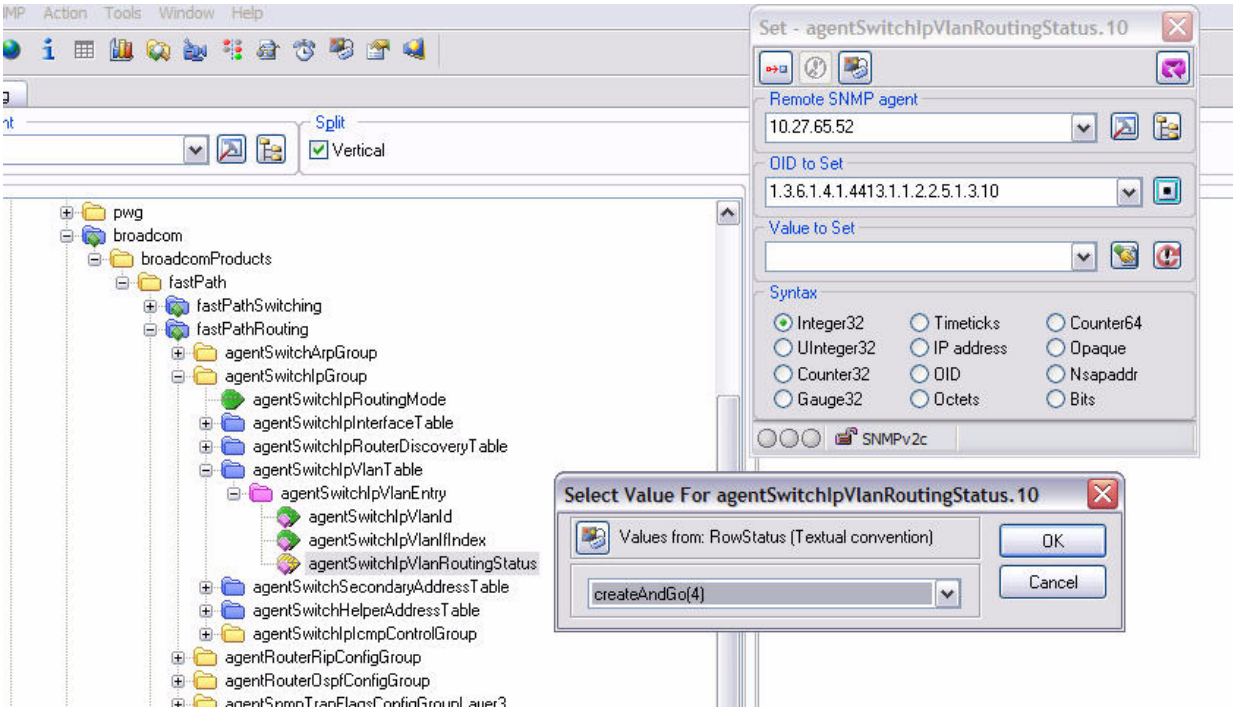

6. Walk the agentSwitchIpVlanIfIndex object to view the logical interface IDs assigned to the VLAN routing interfaces.
7. Set the agentSwitchIpRoutingMode object to enable (1) to enable routing for the switch:
8. Use the agentSwitchIpInterfaceIpAddress and agentSwitchIpInterfaceIpMask objects in the agentSwitchIpInterfaceTable to configure the IP addresses and subnet mask for the virtual router ports.

Note...

While setting the ip address for the VLAN interface, the agentSwitchIpInterfaceIpAddress and agentSwitchIpInterfaceNetMask should be set together.

- VLAN index 482 (VLAN 10): 192.150.3.1 255.255.255.0
- VLAN index 483 (VLAN 20): 192.150.4.1 255.255.255.0

A.4 Configuring OSPF

This section contains two OSPF configuration examples.

Example 1: Configuring an OSPF Border Router and Setting Interface Costs

The following example shows you how to configure an OSPF border router areas and interfaces in FASTPATH 6.2 ENT software.

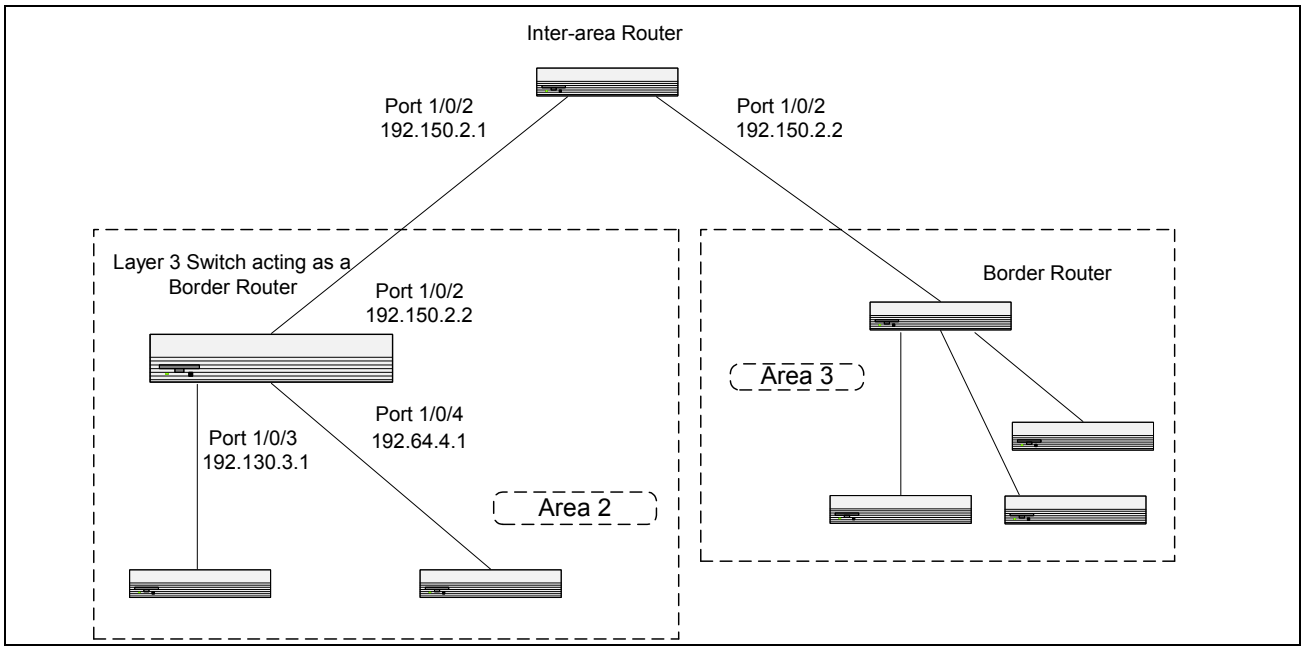

Figure A-3: OSPF Example Network Diagram: Border Router

A.4.1 Using the Web UI to Configure OSPF

1. Go to the **Routing > IP > Configuration** page.
2. From the Routing Mode menu, select Enable to enable routing on the switch.
3. Go to the **Routing > IP > Configuration** page.
4. From the Slot/Port menu, select 0/2.
5. Enter 192.150.2.2 in the IP Address field and 255.255.255.0 in the Subnet Mask field.
6. From the Routing Mode menu, select Enable.

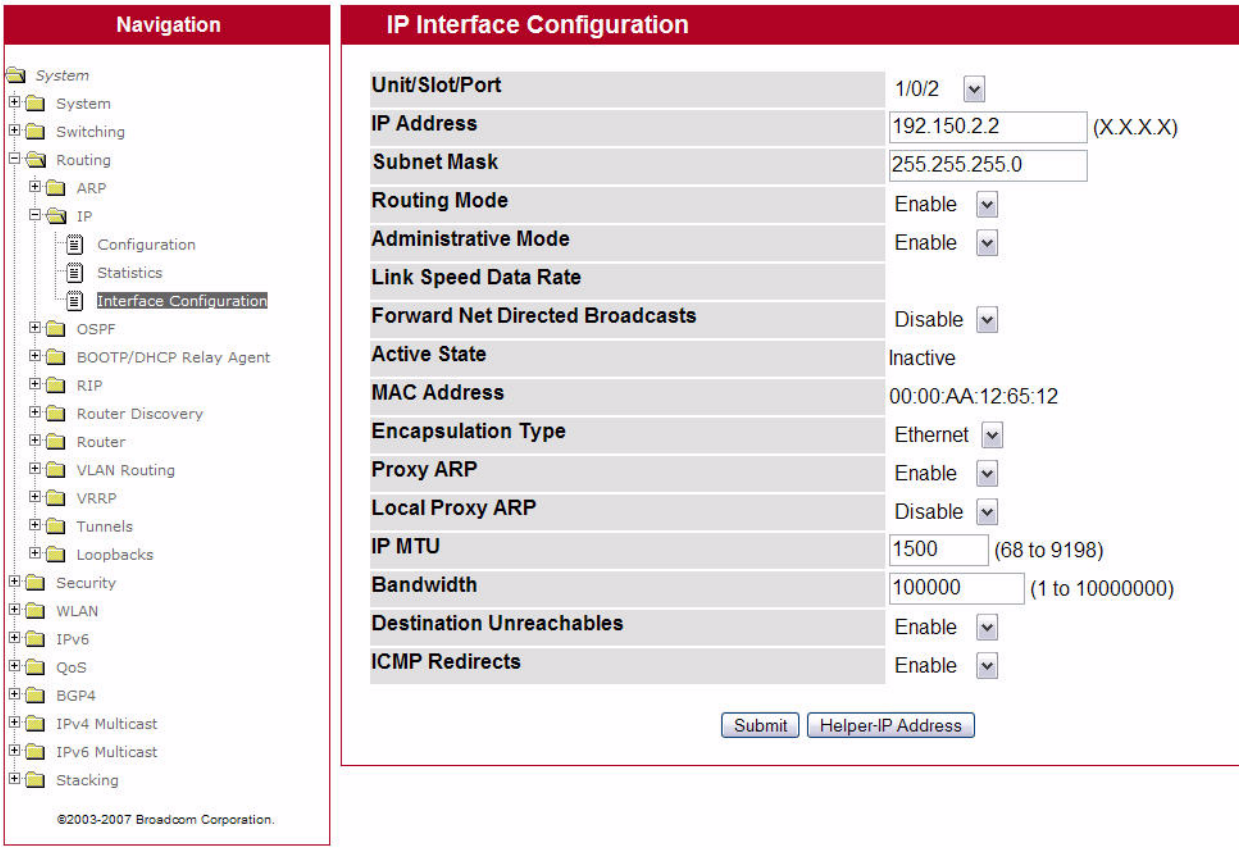

7. Click **Submit**.
8. Configure the IP address and subnet masks for ports 1/0/3 and 1/0/4, and enable routing on both ports.
   - Port 1/0/3: 192.130.3.1/24
   - Port 1/0/4: 192.64.4.1/24
9. Go to the **Routing > OSPF > OSPF Configuration** page.
10. Enter 192.150.9.9 in the Router ID field set the RFC 1583 Compatibility option to Disabled to prevent a routing loop.
11. Click **Submit**.

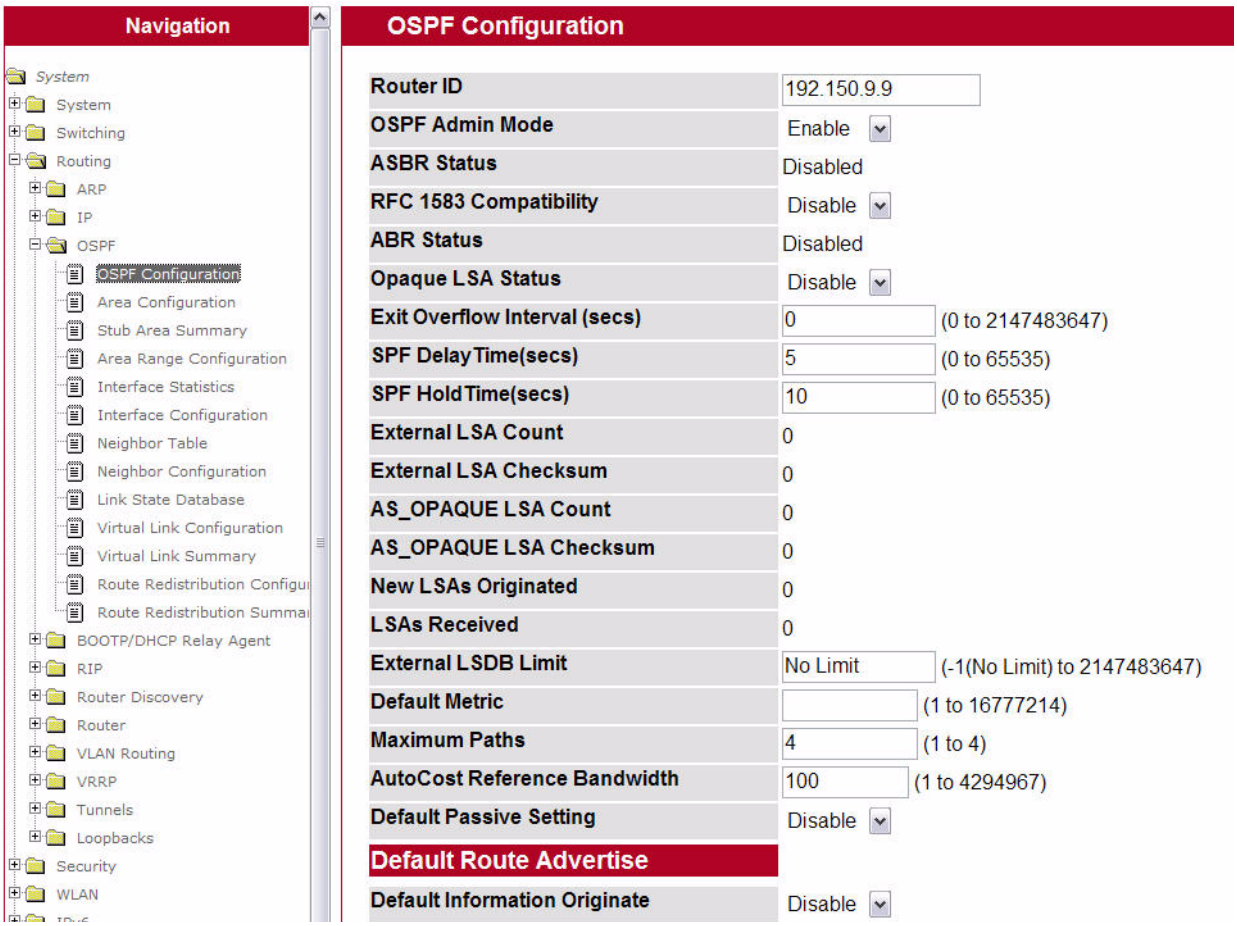

12. Go to the **Routing > OSPF > Interface Configuration** page.

OSPF is globally enabled by default. To make it operational on the router, you configure OSPF for particular interfaces and identify which area the interface is associated with.

13. Select interface 1/0/2 from the Slot/Port menu.

14. From the OSPF Admin Mode field, select Enable.

15. Enter 0.0.0.0 in the OSPF Area ID field.

16. Set the Router Priority field to 128 and the Metric Cost field to 32.

17. Click **Submit.**

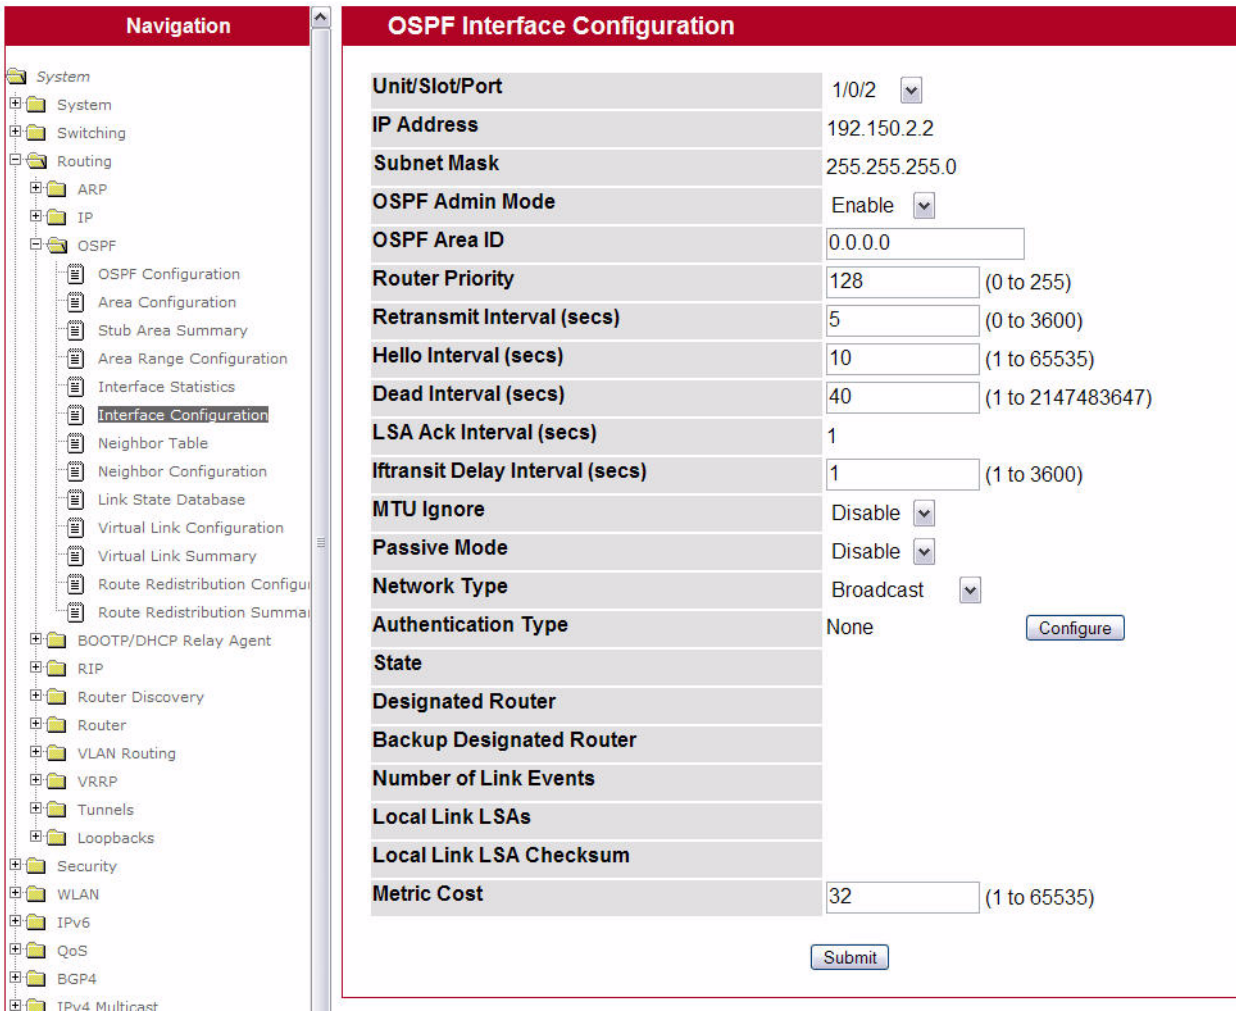

18. Select interface 1/0/2 from the Slot/Port menu.

19. Enable the OSPF Admin Mode and set the Area ID to 0.0.0.3.

20. Set the Router Priority field to 255 and the Metric Cost value to 64.

21. Slick **Submit**.

22. Configure interface 1/0/3 with the following settings, and then click **Submit**:

- OSPF Area ID: 0.0.0.2
- Router Priority 255
- Metric Cost 64

23. Configure interface 1/0/4 with the following settings, and then click **Submit**:

- OSPF Area ID: 0.0.0.2
- Router Priority 255
- Metric Cost 64

A.4.2 Using the CLI to Configure OSPF

1. Enable routing on the switch.

```
(Broadcom FASTPATH Routing) #config
    ip routing
    exit
```

2. For ports 1/0/2, 1/0/3, and 1/0/4, enable routing and assign IP addresses.

```
config
  interface 1/0/2
    routing
    ip address 192.150.2.2 255.255.255.0
    exit
 interface 1/0/3
    routing
    ip address 192.130.3.1 255.255.255.0
    exit
 interface 1/0/4
    routing
    ip address 192.64.4.1 255.255.255.0
    exit
 exit
```

3. Specify a router ID and disable 1583 compatibility to prevent a routing loop (IPv4-only).

```
config
 router ospf
    router-id 192.150.9.9
    no 1583compatibility
    exit
 exit
```

4. Configure the OSPF area ID, priority, and cost for each interface.

   OSPF is globally enabled by default. To make it operational on the router, you configure OSPF for particular interfaces and identify which area the interface is associated with. The following commands also sets the priority and cost for the ports:

```
config
 interface 1/0/2
    ip ospf area 0.0.0.0
    ip ospf priority 128
    ip ospf cost 32
    exit
 interface 1/0/3
    ip ospf area 0.0.0.2
    ip ospf priority 255
    ip ospf cost 64
    exit
 interface 1/0/4
    ip ospf area 0.0.0.2
    ip ospf priority 255
    ip ospf cost 64
    exit
 exit
```

Note...

In OSPFv2, you can also enable OSPF on an interface in global configuration mode by associating a network interface, identified by a network IP address and wild-card mask, with an area. The following example is equivalent to defining interface 1/0/4 in area 2, as in the previous example:

(Broadcom FASTPATH Routing) #config

router ospf

network 192.164.4.0 0.0.0.255 area 2

Example 2: Configuring Stub and NSSA Areas

In this example, Area 0 connects directly to two other areas: Area 1 is defined as a stub area and Area 2 is defined as an NSSA area.

Note...

OSPFv2 and OSPFv3 can operate concurrently on a network and on the same interfaces (although they do not interact). This example configures both protocols simultaneously.

Figure A-4 illustrates this example OSPF configuration.

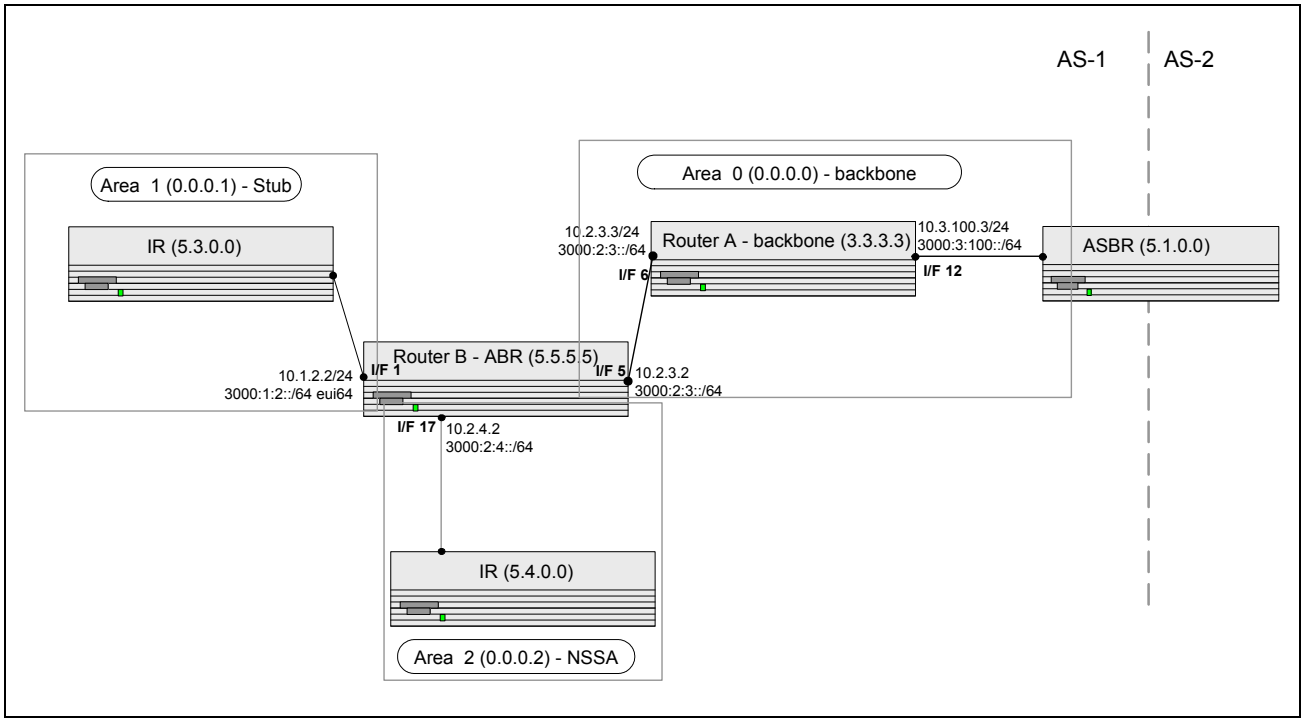

Figure A-4: OSPF Configuration—Stub Area and NSSA Area

www.kontron.com

A.4.3 Using the Web UI to Configure OSPF Areas

Configure Router A: Router A is a backbone router. It links to an ASBR (not defined here) that routes traffic outside the AS.

1. Globally enable IPv6 and IPv4 routing:
   - Use the Routing Mode menu on the **Routing > IP > Configuration** page to enable IPv4 routing.
   - Use the IPv6 Unicast Routing menu on the **IPv6 > Global Configuration** page to enable IPv6 routing.
2. From the **Routing > IP > Configuration** page, enable routing on ports 1/0/6 and 1/0/12 and configure the ports with the following IP addresses:
   - 1/0/6: 10.2.3.3
   - 1/0/12: 10.3.100.3
3. From the **Routing > OSPF > Interface Configuration** page, enable OSPF for ports 1/0/6 and 1/0/12.

   These interfaces are associated with area 0.0.0.0, so you do not need to change the Area ID field.
4. Go to the **IPv6 > Interface Configuration** page and configure the IPv6 interface settings for ports 1/0/6 and 1/0/12.
   a. From the Interface menu, select port 1/0/6.
   b. From the IPv6 Mode menu, select Enable.
   c. In the IPv6 Prefix field, enter 3000:2:3:200/64.
   d. Select the EUI 64 option.
   e. From the Routing Mode field, select Enable.
   f. Click **Submit**.

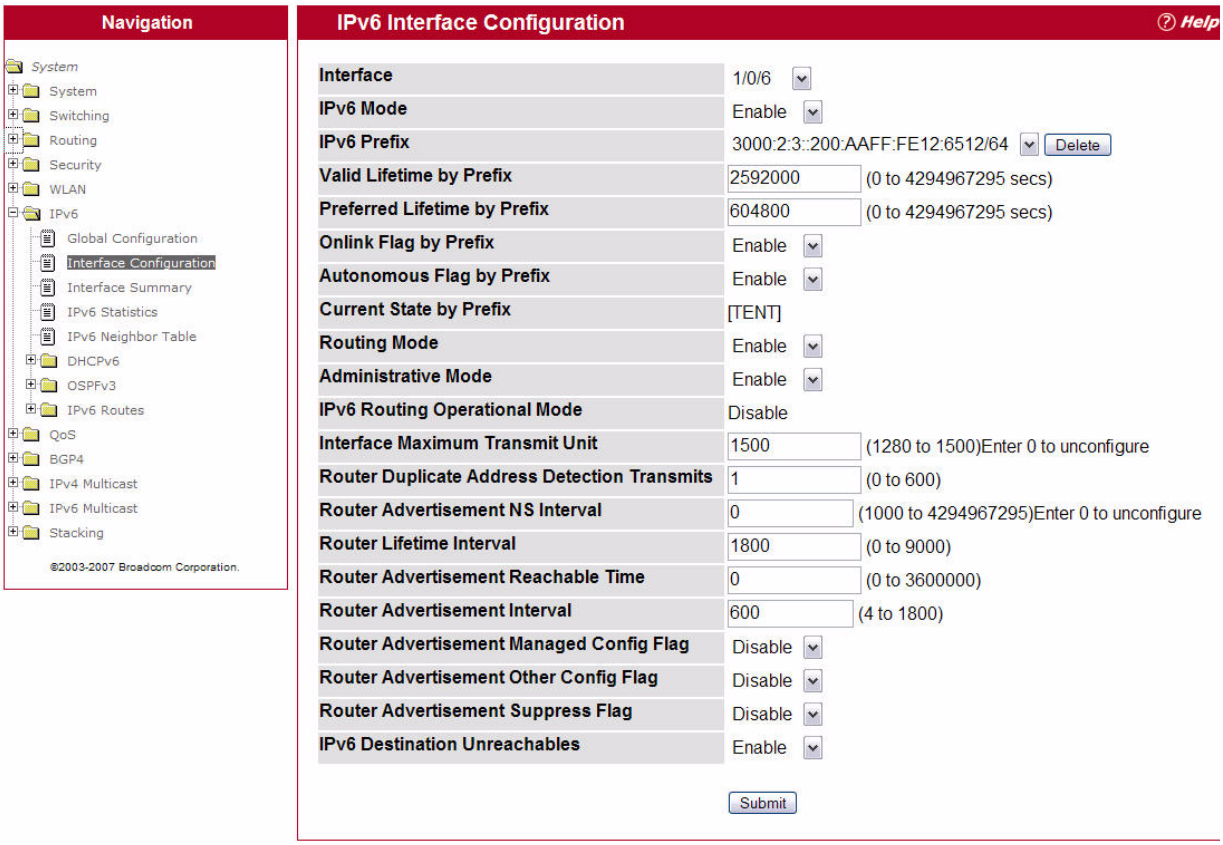

   g. Use the same steps to configure port 1/0/12, but use the IPv6 Prefix 3000:3:100::/64.

5. From the **Routing > OSPF > OSPF Configuration** page, enter 3.3.3.3 in the Router ID field and click **Submit**.

> **Note...**
>
> If the router ID is already configured, you must disable the OSPF Router Admin mode before you change the ID.

Configure Router B: Router B is a ABR that connects Area 0 to Areas 1 and 2.

1. Configure IPv6 and IPv4 routing.

   The static routes are included for illustration only: redistributed static routes, like routes distributed from other protocols, are not injected into stub areas such as Area 1

   Configure a static IPv4 route:

   a. From the **Routing > Router > Configured Routes** page, click **Add Route**.
   b. From the Route Type menu, select Static.
   c. In the Network Address field, enter 10.23.67.0.
   d. In the Subnet Mask field, enter 255.255.255.0.
   e. In the Next Hope Address field, enter 10.2.3.3.

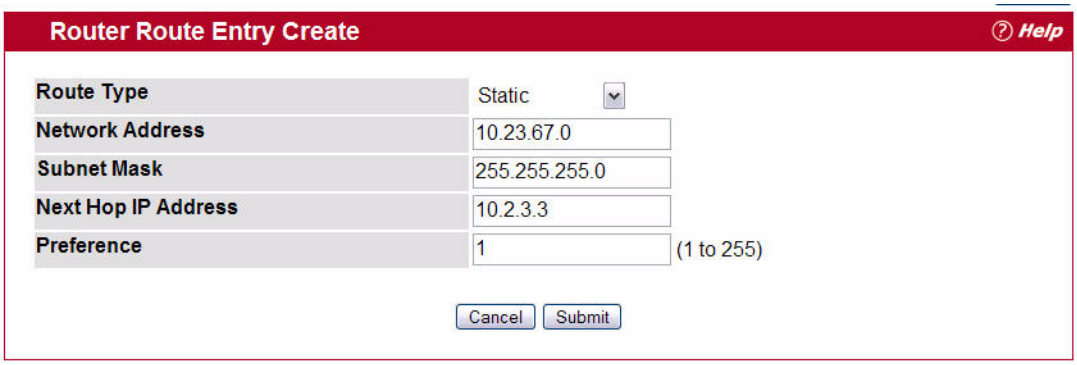

   f. Click **Submit**.

   Configure a static IPv6 route:

   a. Go to the **IPv6 IPv6 Routes > IPv6 Route Entry Configuration** page.
   b. In the IPv6 Network Prefix/Prefix Length field, enter 3000:44:44::/64.
   c. In the Next Hop IPv6 Address field, enter 3000:2:3::210:18ff:fe82:c14.

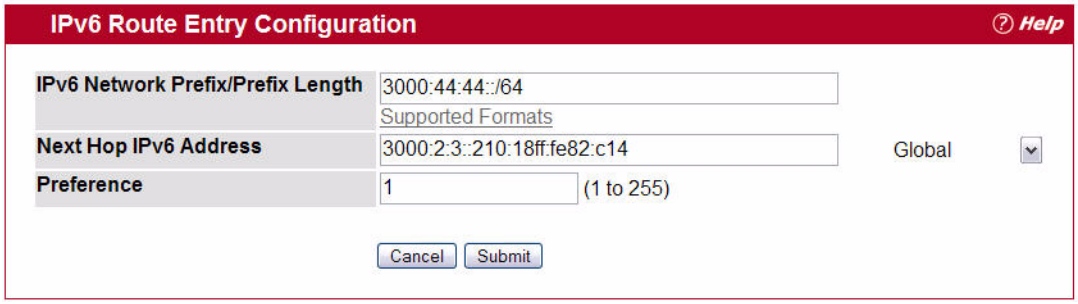

   d. Click **Submit**.

2. From the **Routing > IP > Configuration** page, enable routing on ports 1/0/1, 1/0/5 and 1/0/17 and configure the ports with the following IP addresses:
   - 1/0/1: 10.1.2.2
   - 1/0/5: 10.2.3.2
   - 1/0/17: 10.2.4.2
3. From the **Routing > OSPF > Interface Configuration** page, enable OSPF for ports 1/0/1, 1/0/5, and 1/0/17 and enter the following information in the OSPF Area ID field:
   - 1/0/1: 0.0.0.1
   - 1/0/5: 0.0.0.0
   - 1/0/17: 0.0.02
4. Go to the **IPv6 > Interface Configuration** page and configure the IPv6 interface settings for ports 1/0/1, 1/0/5, and 1/0/17.
   a. From the Interface menu, select port 1/0/1.
   b. From the IPv6 Mode menu, select Enable.
   c. In the IPv6 Prefix field, enter `3000:1:2::/64`.
   d. Select the EUI 64 option.
   e. From the Routing Mode field, select Enable.
   f. Click **Submit**.
   g. Perform the same configuration for interfaces 1/0/5 and 1/0/17 but use the following addresses:
   - 3000:2:3::/64
   - 3000:2:4::/64
5. Go to the **IPv6 > OSPFv3 > Interface Configuration** page enable the OSPFv3 Admin mode and configure the OSPFv3 Area ID for each interface:
   - 1/0/1: Area ID 1
   - 1/0/5: Area ID 0
   - 1/0/17: Area ID 2
6. From the **Routing > OSPF > OSPF Configuration** page, enter 2.2.2.2 in the Router ID field.
7. From the **Routing > OSPF > Area Configuration** page, select area 0.0.0.1 from the Area ID menu.
8. Click **Create Stub Area**.

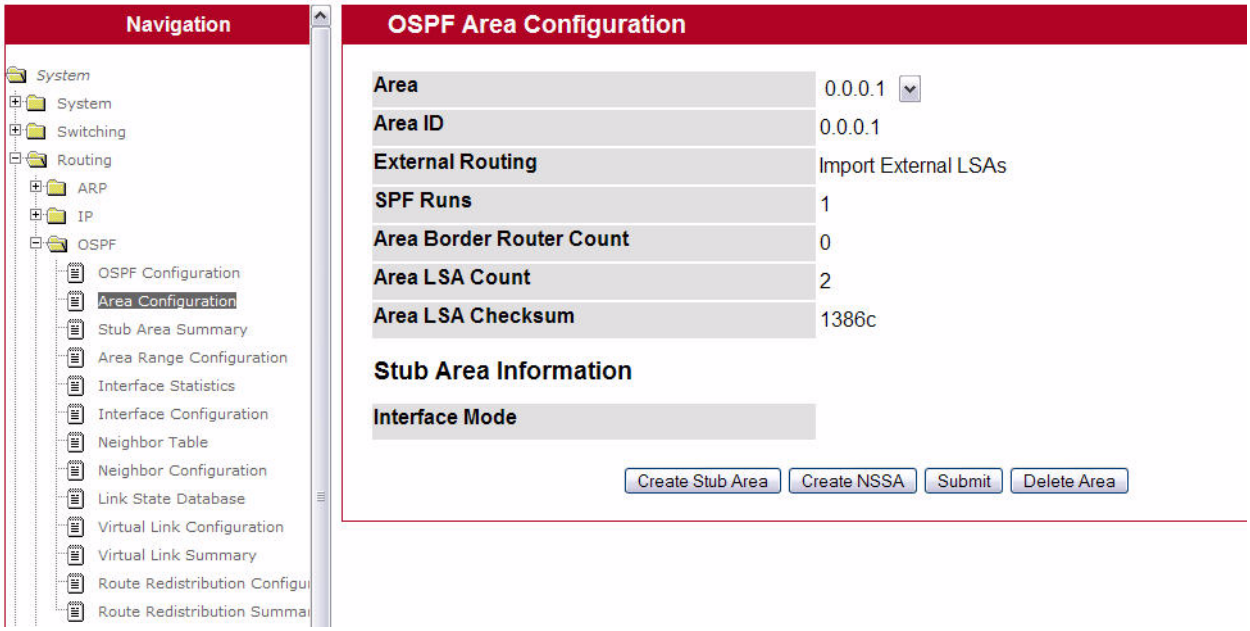

9. Click **Submit** to define Area 1 as a stub.
10. Select area 0.0.0.2 from the Area ID menu and click **Create NSSA**.
11. Click **Submit** to define Area 2 as a Not-So-Stubby-Area (NSSA).
12. To configure a metric cost to associate with static routes when they are redistributed via OSPF, use the following steps.
    a. Go to the **Routing > OSPF > Route Redistribution Configuration** page.
    b. Select Static from the Configured Source Field.
    c. Enter 1 in the Metric field.
    d. Click **Submit**.

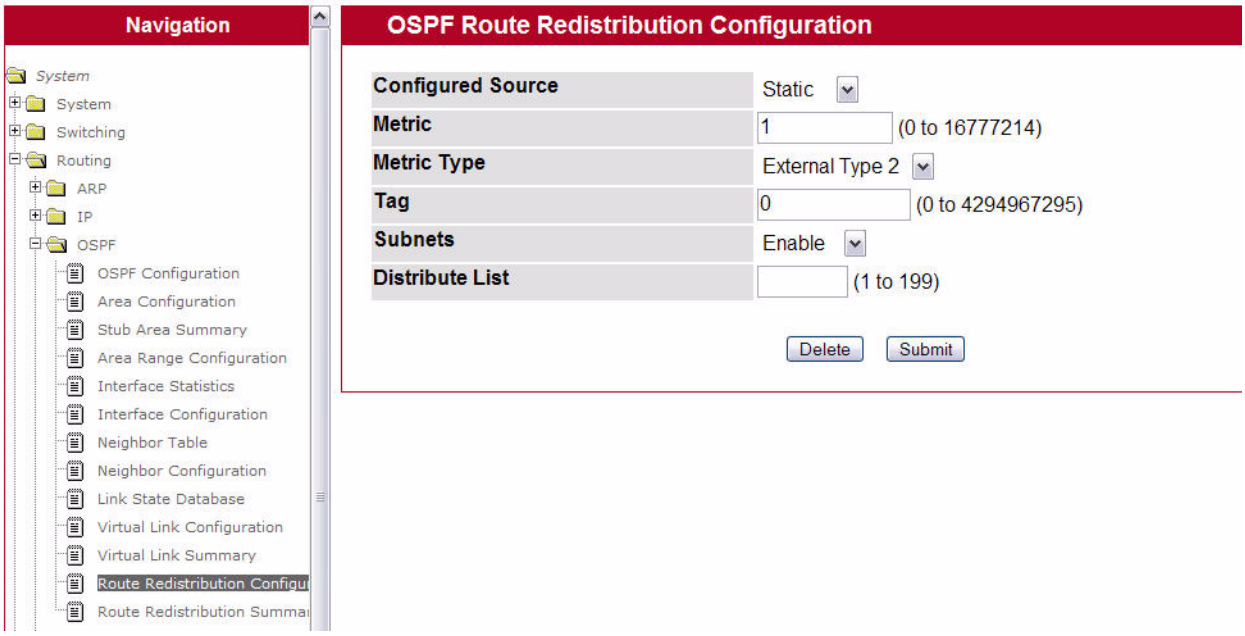

13. Perform the IPv6 OSPFv3 configuration.
    a. From the **IPv6 > OSPFv3 > Configuration** page, enter 2.2.2.2 in the Router ID field to define an OSPF router for IPv6.
14. From the **IPv6 > OSPFv3 > Area Configuration**, define Area 1 as a stub and area 2 as a Not-So-Stubby-Area (NSSA).
15. From the **IPv6 > OSPFv3 > Route Redistribution Configuration** page, configure a metric cost to associate with static routes when they are redistributed via OSPF.

A.4.4 Using the CLI to Configure OSPF Areas

Configure Router A: Router A is a backbone router. It links to an ASBR (not defined here) that routes traffic outside the AS.

1. Globally enable IPv6 and IPv4 routing:
```
Broadcom FASTPATH Routing) #configure
    ipv6 unicast-routing
    ip routing
```
2. Configure IP address and enable OSPF on interfaces 6 and 12 and enable IPv6 OSPF on the interfaces. (OSPF is enabled on the IPv4 interface in the next code group.)
```
    interface 1/0/6
        routing
        ip address 10.2.3.3  255.255.255.0
```

```
        ipv6 address 3000:2:3::/64 eui64
        ipv6 ospf
        exit

    interface 1/0/12
        routing
        ip address  10.3.100.3  255.255.255.0
        ipv6 address 3000:3:100::/64 eui64
        ipv6 ospf
        exit
```

3. Define an OSPF router. Enable OSPF for IPv4 on the two interfaces by globally defining the range of IP addresses associated with each interface, and then associating those ranges with Area 0:

```
    ipv6 router ospf
        router-id 3.3.3.3
        exit
    router ospf
        router-id 3.3.3.3
        network 10.2.3.0 0.0.0.255 area 0.0.0.0
        network 10.3.100.0 0.0.0.255 area 0.0.0.0
        exit
    exit
```

Configure Router B: Router B is a ABR that connects Area 0 to Areas 1 and 2.

1. Configure IPv6 and IPv4 routing. The static routes are included for illustration only: Redistributed static routes, like routes distributed from other protocols, are not injected into stub areas such as Area 1:

```
 Broadcom FASTPATH Routing) #configure
    ipv6 unicast-routing
    ipv6 route 3000:44:44::/64 3000:2:3::210:18ff:fe82:c14
    ip route 10.23.67.0 255.255.255.0 10.2.3.3
```

2. On interfaces 1, 5, and 17, configure IPv4 and IPv6 addresses and enable OSPF on the interfaces. For IPv6, associate interface 1 with Area 1 and interface 17 with Area 2. (OSPF is enabled on the IPv4 interface in the next code group.)

```
    interface 1/0/1
        routing
        ip address  10.1.2.2  255.255.255.0
        ipv6 address 3000:1:2::/64 eui64
        ipv6 ospf
        ipv6 ospf areaid 1
        exit
    interface 1/0/5
        routing
        ip address  10.2.3.2  255.255.255.0
        ipv6 address 3000:2:3::/64 eui64
        ipv6 ospf
        exit
    interface 1/0/17
        routing
        ip address  10.2.4.2  255.255.255.0
        ipv6 address 3000:2:4::/64 eui64
        ipv6 ospf
        ipv6 ospf areaid 2
        exit
```

3. For IPv4: Define an OSPF router. Define Area 1 as a stub. Enable OSPF for IPv4 on interfaces 1, 5, and 17 by globally defining the range of IP addresses associated with each interface, and then associating those ranges with Areas 1, 0, and 17, respectively. Then, configure a metric cost to associate with static routes when they are redistributed via OSPF:

```
    router ospf
        router-id 2.2.2.2
        area 0.0.0.1 stub
        area 0.0.0.2 nssa
```

```
        network 10.1.2.0 0.0.0.255 area 0.0.0.1
        network 10.2.3.0 0.0.0.255 area 0.0.0.0
        network 10.2.4.0 0.0.0.255 area 0.0.0.2
        redistribute static metric 1 subnets
        exit
```

4. For IPv6: Define an OSPF router. Define Area 1 as a stub and area 2 as a Not-So-Stubby-Area (NSSA). Configure a metric cost to associate with static routes when they are redistributed via OSPF:

```
    ipv6 router ospf
        router-id 2.2.2.2
        area 0.0.0.1 stub
        area 0.0.0.2 nssa
        redistribute static metric 105 metric-type 1
        exit
    exit
```

A.5 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 process creates a new authentication list, called radiusList, which uses RADIUS as the authentication method. This authentication list is associated with the 802.1X default login. IEEE 802.1X port-based access control is enabled for the system, and interface 1/0/1 is configured to be in force-authorized mode because this is where the RADIUS server and protected network resources are located.

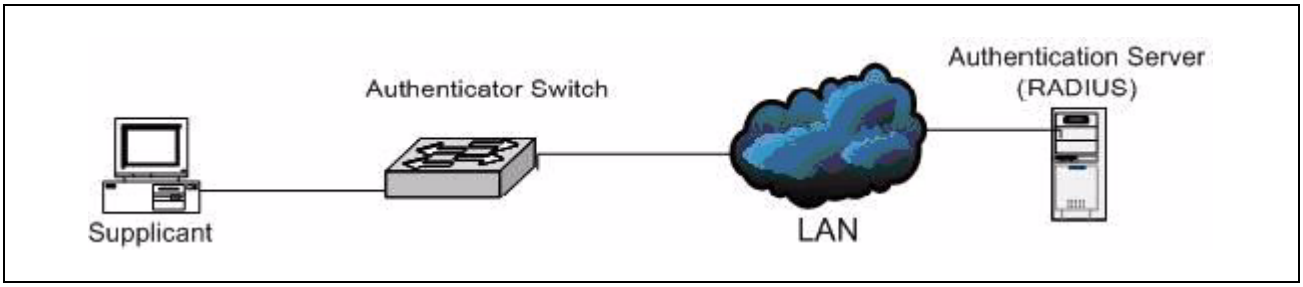

Figure A-5: Switch with 802.1x Network Access Control

If a user, or supplicant, attempts to communicate via the switch on any interface except interface 1/0/1, the system challenges the supplicant for login credentials. The system encrypts the provided information and transmits it to the RADIUS server. If the RADIUS server grants access, the system sets the 802.1X port state of the interface to authorized, and the supplicant is able to access network resources.

A.5.1 Using the Web Interface to configure 802.1X Port-Based Access Control

To configure the RADIUS Server information in the switch, go to the **Security > RADIUS > Server Configuration** page.

Select Add from RADIUS Server Host Address field.

5. Enter 10.10.10.10 in the Host Address field.
6. Click **Submit**.

The page refreshes, and additional fields appear.

7. in the Secret field, enter secret and select the Apply option.
8. From the Primary Server field, select Yes.

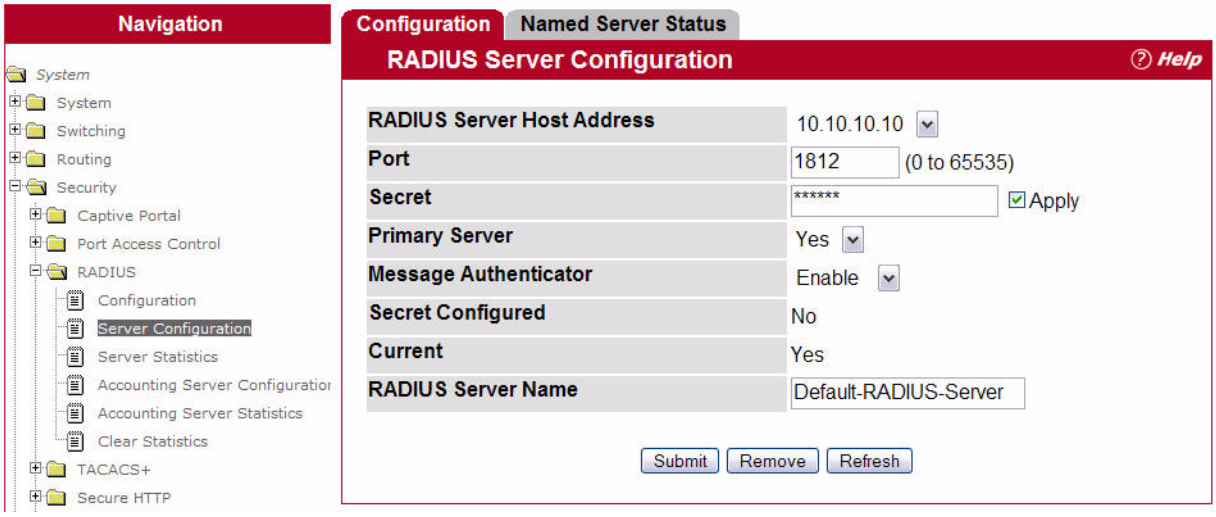

9. Click **Submit** to apply the changes to the system.

Configure the RADIUS accounting server information.

Go to the **Security > RADIUS > Accounting Server** page.

a. Select Add from Accounting Server Host Address field.
b. Enter 10.10.10.10 in the Accounting Server Host Address field.
c. Click **Submit**.
d. in the Secret field, enter secret and select the Apply option.
e. Click **Submit**.

To enable the RADIUS accounting mode, go to the **Security > RADIUS > Configuration** page, select Enable from the Accounting Mode menu, and then click **Submit.**

10. Create an authentication list.
    a. Go to the **System > Configuration > Authentication List Configuration** page.
    b. Enter radiusList in the Authentication List Name field.
    c. Click **Submit**.

Select RADIUS from the Method 1 menu, and then click **Submit.**

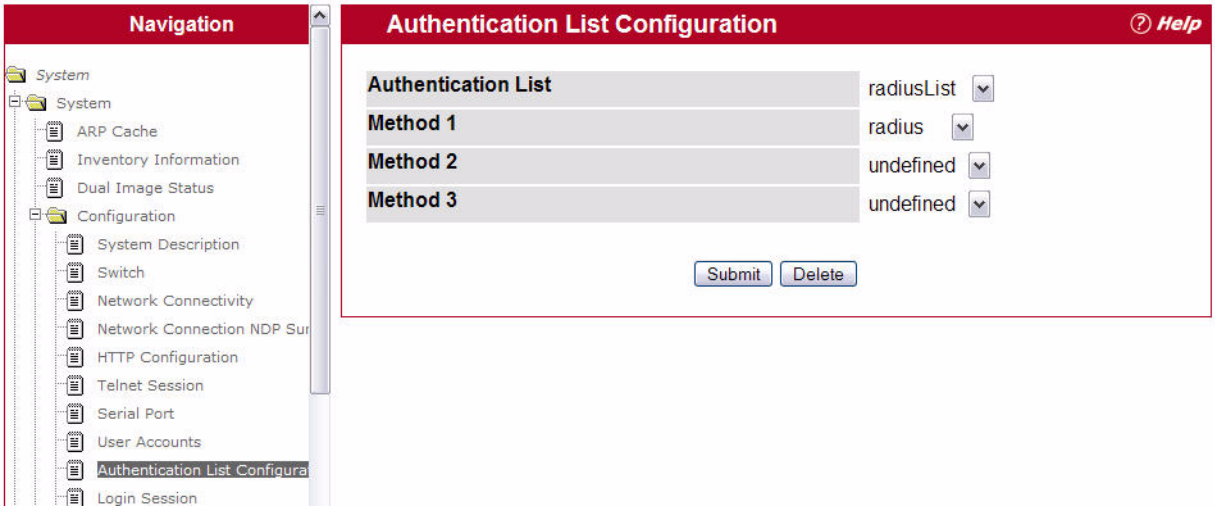

11. To set radiusList as the default login list for users that are not configured on the system, go to the **Switching > Port Access Control > Login** page, select radiusList from the Login field, and click **Submit.**

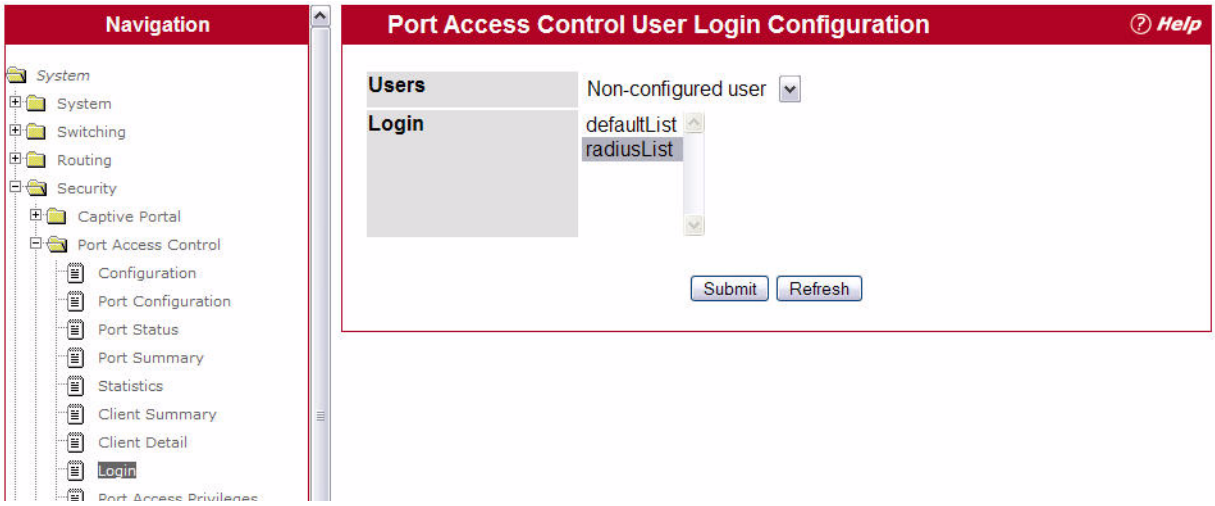

12. To enable IEEE 802.1X authentication on the switch, go to the **Switching > Port Access Control > Configuration** page, select Enable from the Administrative Mode menu, and then click **Submit**.

13. To set the 802.1X mode for port 1/0/1, go to the **Switching > Port Access Control > Port Configuration** page, select Force Authorized from the Control Mode field, and then click **Submit**.

A.5.2 Using the CLI to configure 802.1X Port-Based Access Control

Configure the RADIUS authentication server IP address.

```
(Broadcom FASTPATH Routing) #config
      radius server host auth 10.10.10.10
```

Configure the RADIUS authentication server secret.

```
radius server key auth 10.10.10.10
   secret
   secret
```

Configure the RADIUS accounting server IP address.

```
radius server host acct 10.10.10.10
```

Configure the RADIUS accounting server secret.

```
radius server key acct 10.10.10.10
   secret
   secret
```

Enable RADIUS accounting mode.

```
radius accounting mode
```

14. Create an authentication list named radiusList and set radius as the login method.

```
authentication login radiusList radius
```

15. Set radiusList as the default login list for users that are not configured on the system

```
dot1x default-login radiusList
```

16. Enable 802.1X authentication on the switch.

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

17. Set the 802.1X mode for port 1/0/1 to Force Authorized.

```
interface 1/0/1
   dot1x port-control force-authorized
   exit
```

A.5.3 Using SNMP to configure 802.1X Port-Based Access Control

Use the agentRadiusServerStatus in the agentRadiusServerConfigTable under the FASTPATH-RADIUS-AUTH-CLIENT-MIB to create a new RADIUS server entry.

Use the agentRadiusServerAddress object to configure the RADIUS authentication server IP address as 10.10.10.10.

Use the agentRadiusServerSecret object to configure the RADIUS authentication server secret.

Use the agentRadiusAccountingStatus object in the agentRadiusAccountingConfigTable to create a RADIUS accounting server.

User the agentRadiusAccountingServerAddress object to configure the RADIUS accounting server IP address. as 10.10.10.10.

Use the agentRadiusAccountingSecret object to configure the RADIUS accounting server secret.

Use the agentRadiusAccountingStatus object to enable RADIUS accounting mode.

18. Use the agentAuthenticationListCreate object under the agentAuthenticationGroup in the FASTPATH-SWITCHING-MIB to create an authentication list.

Set the value of the agentAuthenticationListCreate object to radiusList

Set the value of the agentAuthenticationListMethod1 object in the agentAuthenticationListTable to radius (2) to use RADIUS as the login method.

19. Use the agentUserConfigDefaultAuthenticationList object in agentAuthenticationGroup in the FASTPATH-SWITCHING module to set radiusList as the default login list for users that are not configured on the system.

20. To enable 802.1X authentication on the switch, set the dot1xPaeSystemAuthControl object in the IEEE8021-PAE-MIB module to enable (1).

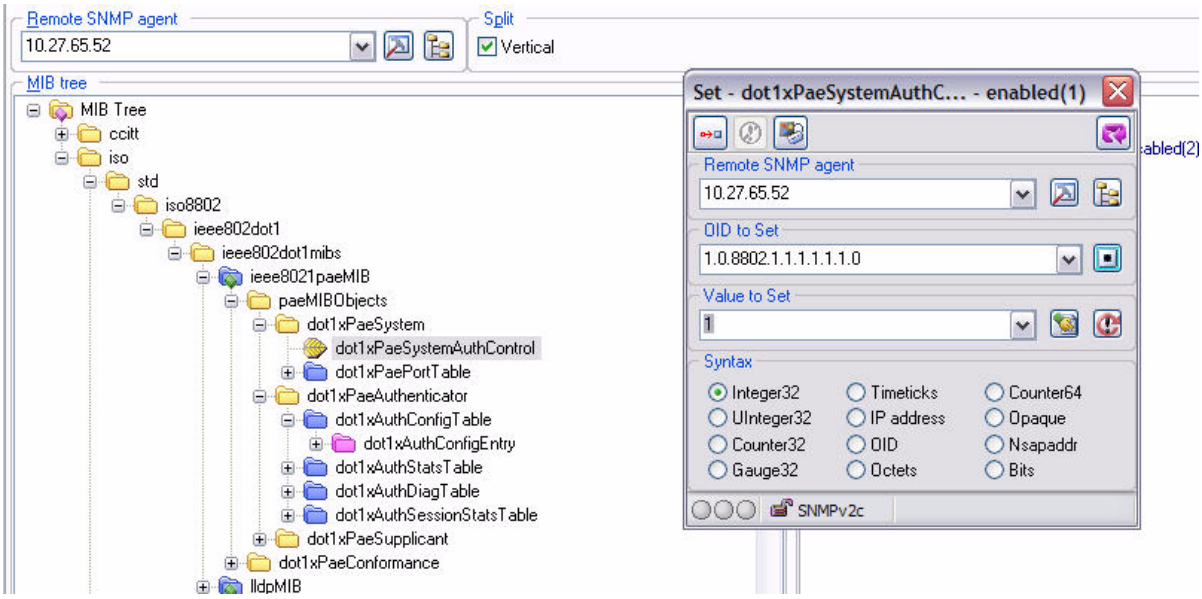

To set the 802.1X mode for port 1/0/1 to Force Authorized, use theagentDot1xPortControlMode object in the agentDot1xPortConfigTable, which is in FASTPATH-DOT1X-ADVANCED-FEATURES-MIB.

A.6 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.

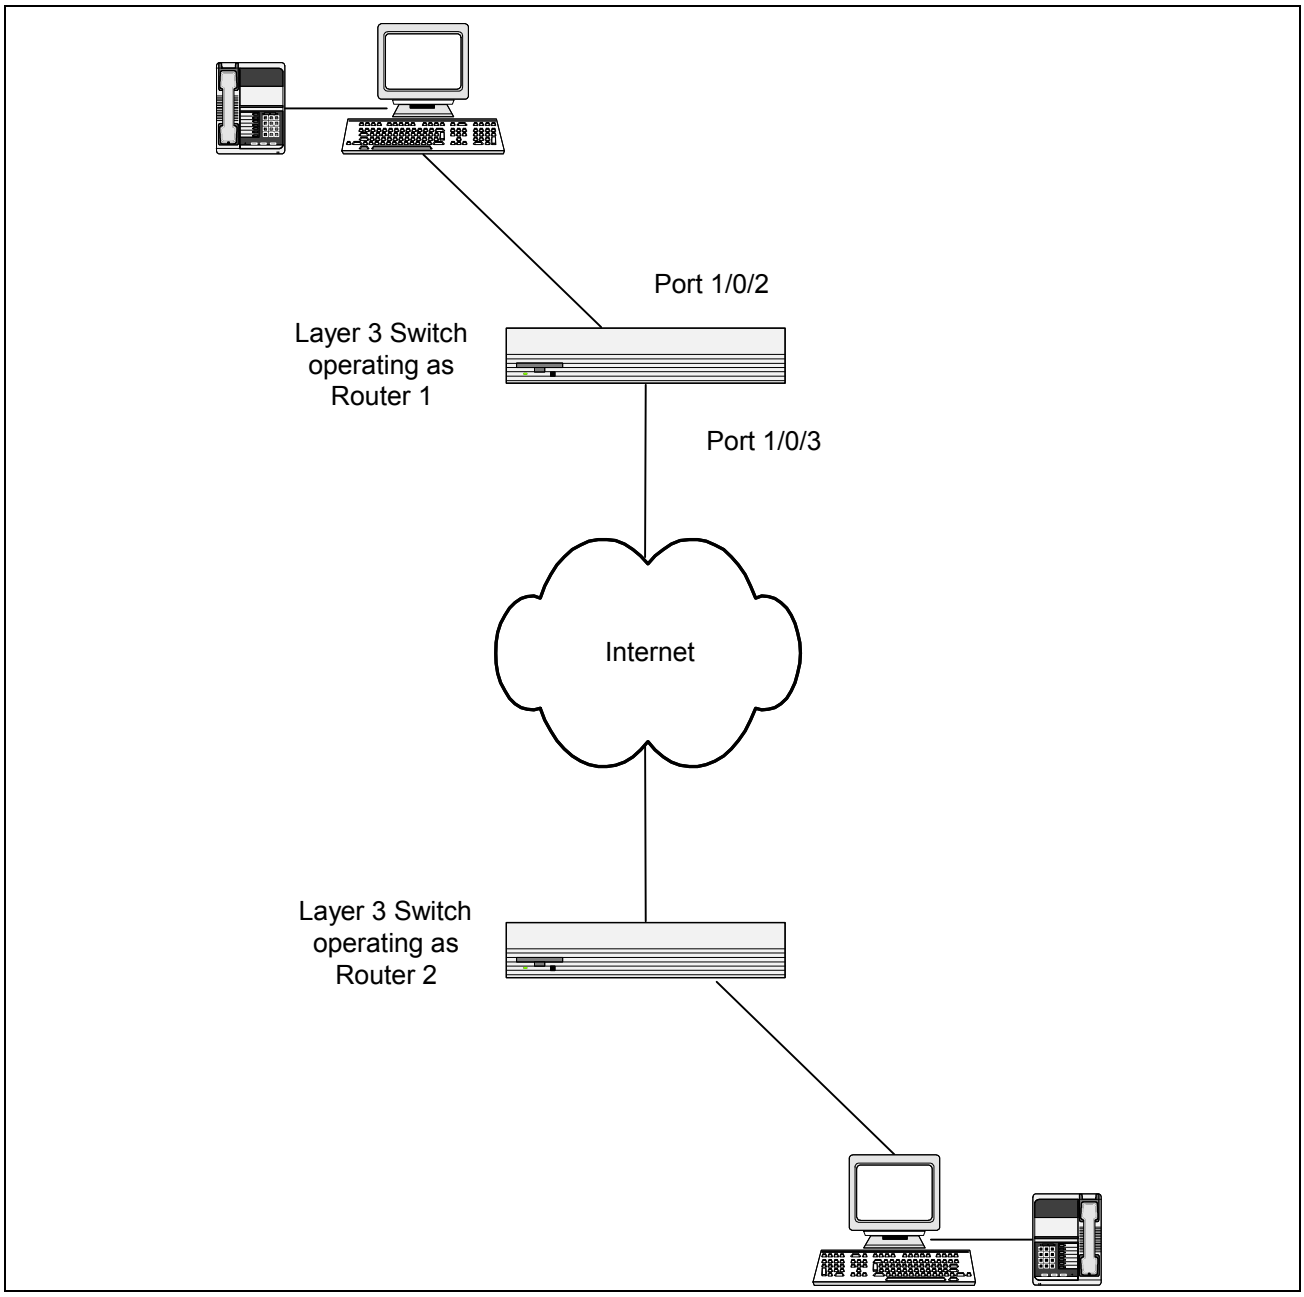

Figure A-6: DiffServ VoIP Example Network Diagram

A.6.1 Using the Web UI to Configure DiffServ VoIP Support

1. To set queue 5 on all ports to use strict priority mode, go to the **QoS > Class of Service > Interface Queue Configuration** page and configure the following settings:
   - Slot/Port: Global
   - Queue ID: 5
   - Scheduler Type: Strict
2. Click **Submit**.

   Queue 5 will be used for all VoIP packets.

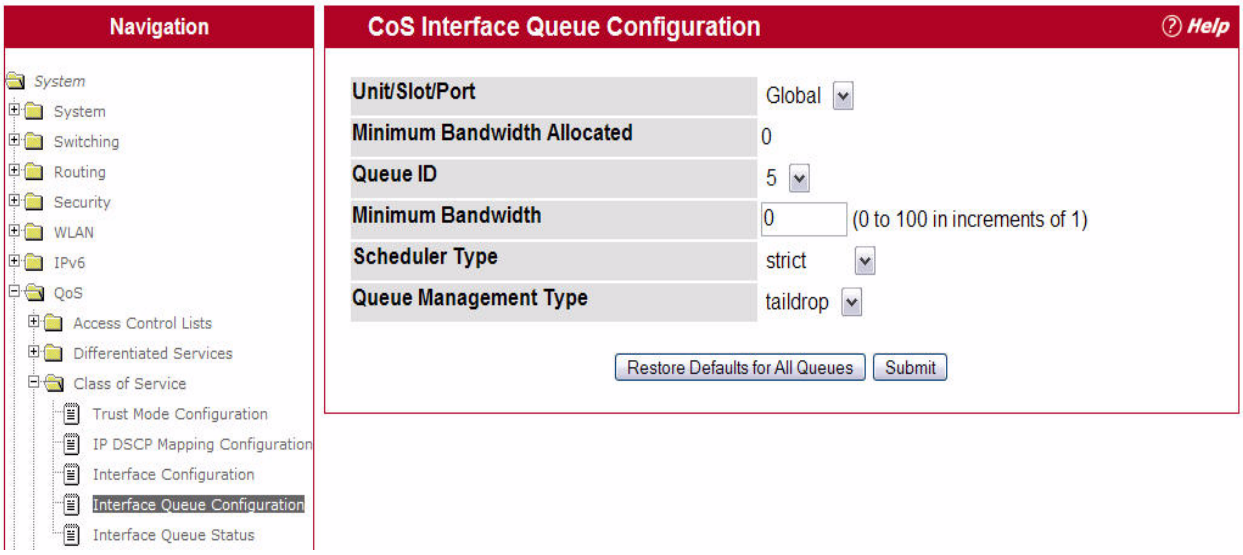

3. Go to the **QoS > Differentiated Services > DiffServ Configuration** page and enable DiffServ for the switch.
4. Go to the **Class Configuration** page, select Create from the Class Selector field, enter class_voip in the Class Name field, select All as the Class Type, and then click **Submit**.
5. Select IPv4 as the Class Layer 3 Protocol, and then click **Submit**.
6. Select Protocol from the Class Match Selector menu, and then click **Add Match Criteria**.
7. Select UDP from the Protocol Keyword menu, and then click **Submit**.

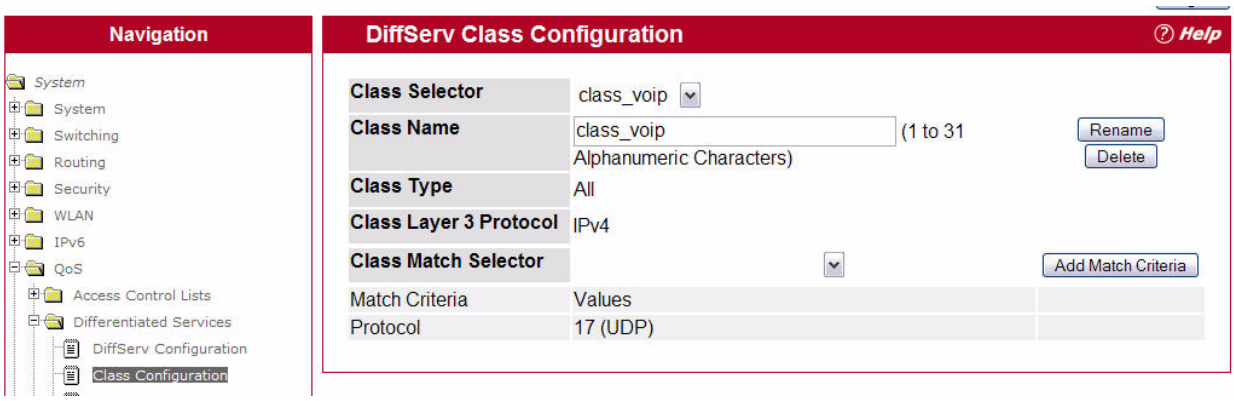

8. Create a second DiffServ classifier named class_ef and define a single match criterion to detect a DiffServ code point (DSCP) of ef (expedited forwarding).

   This handles incoming traffic that was previously marked as expedited elsewhere in the network.

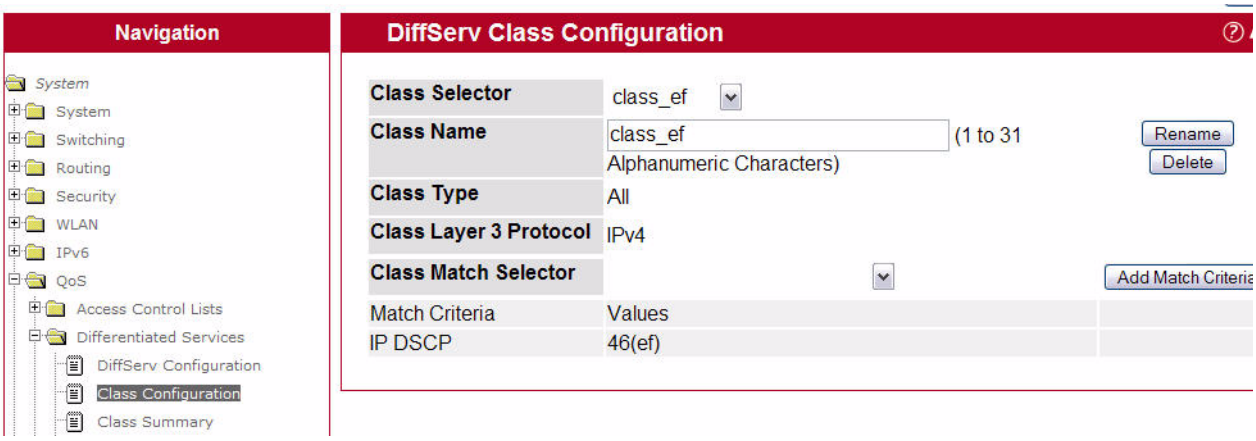

9. Go to the **Policy Configuration** page, select Create from the Policy Selector menu, enter pol_voip in the Policy Name field, and then click **Submit.**
10. From the Available Class List menu, select class_voip, and then click **Add Selected Class**.
11. From the Available Class List menu, select class_ef, and then click **Add Selected Class**.

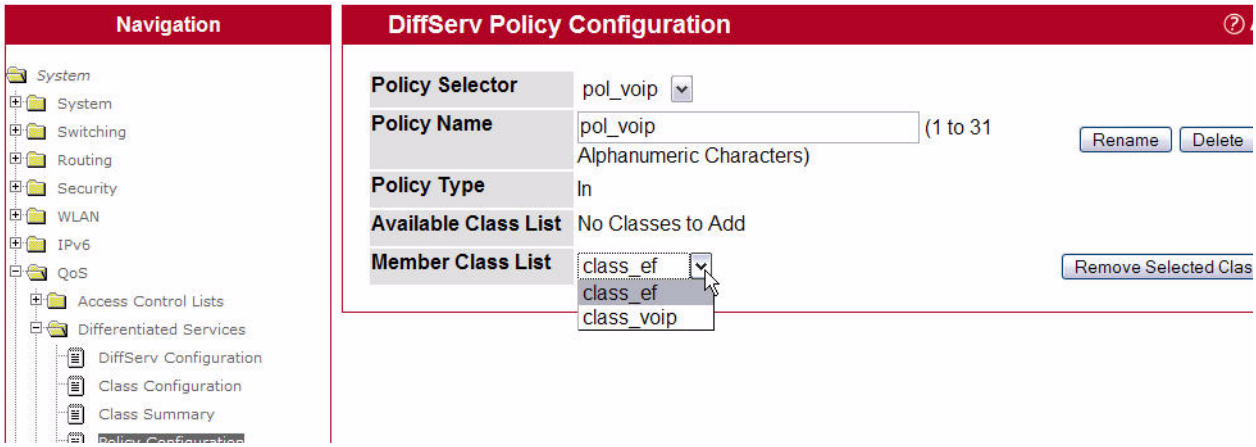

12. Go to the **Policy Class Definition** page and configure how classes that match the policy are handled.

    The following steps configure this policy so that incoming packets already marked with a DSCP value of “EF” (per the class_ef definition), or marks UDP packets (per the class_voip definition) with a DSCP value of “EF.” In both cases, the matching packets are assigned internally to use queue 5 of the egress port to which they are forwarded.

    a. Select pol_viop from the Policy Selector menu, class_ef from the Member Class List menu, and Assign Queue from the Policy Attribute Selector, and then click **Configure Selected Attribute**.
    b. In the Queue ID Value field, enter 5, and then click **Submit**.
    c. Select pol_viop from the Policy Selector menu, class_voip from the Member Class List menu, and Assign Queue from the Policy Attribute Selector, and then click **Configure Selected Attribute**.
    d. Select ef from the DSCP Keyword menu, and then click **Submit**.
    e. Select pol_viop from the Policy Selector menu, class_voip from the Member Class List menu, and Mark IP DSCP from the Policy Attribute Selector, and then click **Configure Selected Attribute**.
    f. Select ef from the DSCP Keyword menu, and then click **Submit**.

13. To attach the defined policy to an inbound service interface, go to the **Service Configuration** page.
14. Select interface 0/2 from the Slot/Port menu.
15. Select pol_voip from the Policy In menu.
16. Click **Submit**.

A.6.2 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.
```
(Broadcom FASTPATH Routing) #config
   cos-queue strict 5
   diffserv
```
2. Create a DiffServ classifier named 'class_voip' and define a single match criterion to detect UDP packets. The class type match-all indicates that all match criteria defined for the class must be satisfied in order for a packet to be considered a match.
```
   class-map match-all class_voip
      match protocol udp
      exit
```
3. Create a second DiffServ classifier named 'class_ef' and define a single match criterion to detect a DiffServ code point (DSCP) of 'EF' (expedited forwarding). This handles incoming traffic that was previously marked as expedited elsewhere in the network.
```
   class-map match-all class_ef
      match ip dscp ef
      exit
```
4. Create a DiffServ policy for inbound traffic named 'pol_voip', then add the previously created classes 'class_ef' and 'class_voip' as instances within this policy.

   This policy handles incoming packets already marked with a DSCP value of 'EF' (per 'class_ef' definition), or marks UDP packets per the 'class_voip' definition) with a DSCP value of 'EF'. In each case, the matching packets are assigned internally to use queue 5 of the egress port to which they are forwarded.
```
   policy-map pol_voip in
      class class_ef
      assign-queue 5
      exit
   class class_voip
      mark ip-dscp ef
      assign-queue 5
      exit
   exit
```
   Attach the defined policy to an inbound service interface.
```
   interface 1/0/2
      service-policy in pol_voip
      exit
   exit
```

A.6.3 Using SNMP to Configure DiffServ VoIP Support

1. Use the agentDiffServGenStatusAdminMode object in agentDiffServGenStatusGroup under fastPathQOSDiffServPrivate in the FASTPATH-QOS-DIFFSERV-PRIVATE-MIB module to activate DiffServ for the switch.
2. To set queue 5 on all ports to use strict priority mode, use the agentCosQueueSchedulerType in the agentCosQueueTable in the FASTPATH-QOS-COS-MIB module. This queue shall be used for all VoIP packets.
3. Use the agentDiffServClassRowStatus object in the agentDiffServClassTable to create two new DiffServ instances. Se the value to CreateAndGo (4).

4. Use the agentDiffServClassName in the agentDiffServClassTable to name the first DiffServ classifier “class_voip” and the second classifier “class_ef.”
5. Use the agentDiffServClassType in the agentDiffServClassTable to set the class type for each classifier to All (1).
6. Use the agentDiffServClassRuleMatchEntryType in the agentDiffServClassRuleTable to set class_voip to match a protocol (9) and class_ef to match an IP DSCP value (6).
7. For class_voip, define a single match criterion to detect UDP packets by setting the agentDiffServClassRuleMatchProtocolNum in the agentDiffServClassRuleTable to 17.
8. Use the agentDiffServClassRuleMatchIpDscp object in the agentDiffServClassRuleTable to define a single match criterion to detect a DSCP of EF (46). This handles incoming traffic that was previously marked as expedited elsewhere in the network.
9. Use the agentDiffServPolicyRowStatus object in the agentDiffServPolicyTable to create a DiffServ policy. Set the value to CreateAndGo (4).
10. Use the agentDiffServPolicyType object to set the policy direction so that it applies to inbound (1) traffic.
11. Use the agentDiffServPolicyName object to name the new DiffServ instance “pol_voip.”
12. Use the agentDiffServPolicyInstRowStatus object in the agentDiffServPolicyInstTable to create new instances that will be associated with the previously created classes (class_ef and class_voip).
13. Use the agentDiffServPolicyInstClassIndex object to associate class_ef and class_voip with the policy instances.
14. Use the agentDiffServPolicyAttrRowStatus object in the agentDiffServPolicyAttrTable to create three instances.
15. Use the agentDiffServPolicyAttrStmtAssignQueueId to set the queue value for instances 1.1.1 and 1.2.2 to 5. so that matching packets are assigned internally to use queue 5 of the egress port to which they are forwarded.
16. Use the agentDiffServPolicyAttrStmtMarkIpDscpVal object to set the value of instance 1.2.1 to 46, which marks UDP packets (per the class_voip definition) with a DSCP value of EF.
17. Create an instance for the interface that will have the policy attached by using the agentDiffServServiceRowStatus object in the agentDiffServServiceTable. For example, to create an instance for interface 1/0/2, set 2.1 to CreateAndG0 (4).
18. Attach the policy to the interface instance by using the agentDiffServServicePolicyIndex object. Set the value of the instance to 1.

A.7 Configuring PIM

Protocol Independent Multicast (PIM) is a standard multicast routing protocol that provides scalable inter-domain multicast routing across the Internet, independent of the mechanisms provided by any particular unicast routing protocol.

PIM-SM is used to efficiently route multicast traffic to multicast groups that may span wide area networks where bandwidth is a constraint. PIM-SM is defined in RFC 4601.

The following example configures PIM-SM for IPv4 on a router.

A.7.1 Using the Web UI to Configure PIM-SMv4

1. Access the **Routing > OSPF > OSPF Configuration** page and enter 3.3.1.1 in the Router ID field, and then click **Submit** to configure an OSPF[1] router.
2. Globally enable IP routing, multicast, IGMP, and PIM-SM.
   - For IP routing: go to the **Routing > IP > Configuration** page and select Enable in the Routing Mode field.

1. OSPF configuration is added as a unicast protocol for illustration purposes; static unicast routing could also be configured.

- For Multicast, go to the **IPv4 Multicast > Global Configuration** page and select Enable from the Admin Mode field.
- For IGMP, go to the **IPv4 Multicast > IGMP** > **Global Configuration** page and select Enable from the Admin Mode field.
- For PIM-SIM, go to the **IPv4 Multicast > PIM-SM** > **Global Configuration** page and select Enable from the Admin Mode field.

3. From the **IPv4 Multicast > PIM-SM > Static RP Configuration** page, configure a PIM-SM rendezvous point with an IP address and group range. The IP address will serve as an RP for the range of potential multicast groups specified in the group range.
   a. In the RP Address field, enter 1.1.1.1.
   b. In the Group Address field, enter 224.0.0.0.
   c. In the Group Mask field, enter 224.0.0.0.

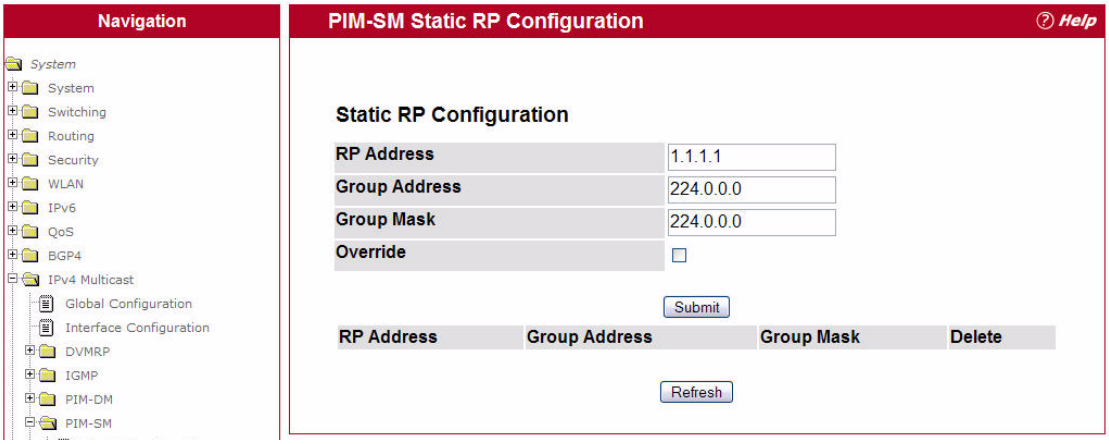

   d. Click **Submit**.

4. Enable routing, IGMP, PIM-SM, and OSPF on one or more interfaces by going to the pages listed below, selecting the interface from the Slot/Port menu, enabling the feature, and clicking **Submit**.
   - To enable IP routing on an interface, go to the **Routing > IP > Interface Configuration** page, enter an IP address and Subnet mask into the appropriate fields, and select Enable in the Routing Mode field.
   - For IGMP, go to the **IPv4 Multicast > IGMP** > **Routing Interface > Interface Configuration** page and select Enable from the Interface Mode field.
   - For OSPF, go to the **Routing > OSPF > Interface Configuration** page and select Enable from the OSPF Admin Mode field.
   - For PIM-SIM, go to the **IPv4 Multicast > PIM-SM** > **Interface Configuration** page and select Enable from the Admin Mode field.

A.7.2 Using the CLI to Configure PIM-SMv4

1. Configure an OSPF router and globally enable IP routing, multicast, IGMP, and PIM-SM.

```
(Broadcom FASTPATH Routing) #configure
   router ospf
      router-id 3.3.1.1
      exit
   ip routing
   ip multicast
   ip igmp
   ip pimsm
```

Note...

This router should be an RP.

2. Configure a PIM-SM rendezvous point with an IP address and group range. The IP address will serve as an RP for the range of potential multicast groups specified in the group range.
```
ip pimsm rp-address 1.1.1.1 224.0.0.0 240.0.0.0
```
3. Enable routing, IGMP, PIM-SM, and OSPF on one or more interfaces.
```
interface 1/0/1
   routing
   ip address 3.3.3.1 255.255.255.0
   ip pimsm
   ip igmp
   ip ospf area 0
   exit
interface 1/0/3
   routing
   ip address 1.1.1.1 255.255.255.0
   ip pimsm
   ip igmp
   ip ospf area 0
   exit
exit
```

A.7.3 Using SNMP to Configure PIM-SMv4

1. Use the following objects to configure an OSPF router and globally enable IP routing, multicast, IGMP, and PIM-SM.
   - Enable OSPF: ospfAdminStat under ospfGeneralGroup in the OSPF-MIB module
   - Set OSPF router ID: ospfRouterId under ospfGeneralGroup in the OSPF-MIB module
   - Enable routing: agentSwitchIpRoutingMode object in agentSwitchIpGroup under fastPathRouting
   - Enable multicast: agentMulticastRoutingAdminMode under agentMulticastRoutingConfigGroup in the FASTPATH-MULTICAST-MIB module
   - Enable IGMP: agentMulticastIGMPAdminMode under agentMulticastIGMPConfigGroup
   - Enable PIM-SM: agentMulticastPIMSMAdminMode under agentMulticastPIMSMConfigGroup
2. Use the pimSmStaticRPIPAddress object in the agentMulticastPIMSMStaticRPTable under agentMulticastPIMSMConfigGroup to configure a PIM-SM rendezvous point with an IP address (1.1.1.1) and group range 224.0.0.0 to 240.0.0.0. The IP address will serve as an RP for the range of potential multicast groups specified in the group range.
3. Use the following objects to enable routing, IGMP, PIM-SM, and OSPF on one or more interfaces:
   - Enable routing on the interface: agentSwitchIpInterfaceRoutingMode in the agentSwitchIpInterfaceTable under the FASTPATH-ROUTING-MIB module.
   - Enable IGMP on the interface: mgmdRouterInterfaceStatus in the mgmdRouterInterfaceTable under the MGMD-STD-MIB module.
   - Enable PIM-SM on an interface: pimSmInterfaceStatus in the pimSmInterfaceTable under the PIM-STD-MIB module.
   - Enable OSPF on an interface: ospfIfStatus in the ospfIfTable in the OSPF-MIB module.
4. Use the agentSwitchIpInterfaceIpAddress and agentSwitchIpInterfaceNetMask objects in the agentSwitchIpInterfaceTable under FASTPATH-ROUTING-MIB to assign an IP address and subnet mask to each interface.