ExtremeWare 72.e Installation and User Guide

Software Version 7.2e

Extreme Networks, Inc.
3585 Monroe Street
Santa Clara, California 95051
(888) 257-3000
http://www.extremenetworks.com

Published: February 12, 2004
Part number: 100157-00 Rev 01

Authors: Hugh Bussell, Julie Laccabue, Megan Mahar, Richard Small

Production: Hugh Bussell

Contents

Chapter 15 Using ExtremeWare Vista on the Summit 400

Appendix A Technical Specifications

Appendix B Software Upgrade and Boot Options

Appendix C Troubleshooting

Preface

This preface provides an overview of this guide, describes guide conventions, and lists other publications that might be useful.

Introduction

This guide provides the required information to install the Summit 400-48 switch and configure the ExtremeWare™ software running on the Summit 400-48 switch.

This guide is intended for use by network administrators who are responsible for installing and setting up network equipment. It assumes a basic working knowledge of:

- Local area networks (LANs)
- Ethernet concepts
- Ethernet switching and bridging concepts
- Routing concepts
- Internet Protocol (IP) concepts
- Routing Information Protocol (RIP) and Open Shortest Path First (OSPF).
- IP Multicast concepts
- Protocol Independent Multicast (PIM) concepts
- Simple Network Management Protocol (SNMP)

NOTE

If the information in the release notes shipped with your switch differs from the information in this guide, follow the release notes.

Conventions

Table 1 and Table 2 list conventions that are used throughout this guide.

Table 1: Notice Icons

Icon	Notice Type	Alerts you to...
	Note	Important features or instructions.
	Caution	Risk of personal injury, system damage, or loss of data.
	Warning	Risk of severe personal injury.

Table 2: Text Conventions

Convention	Description
`Screen displays`	This typeface indicates command syntax, or represents information as it appears on the screen.
The words “enter” and “type”	When you see the word “enter” in this guide, you must type something, and then press the Return or Enter key. Do not press the Return or Enter key when an instruction simply says “type.”
[Key] names	Key names are written with brackets, such as [Return] or [Esc].
	If you must press two or more keys simultaneously, the key names are linked with a plus sign (+). Example:
	Press [Ctrl]+[Alt]+[Del].
Words in *italicized* type	Italics emphasize a point or denote new terms at the place where they are defined in the text.

Related Publications

The publications related to this one are:

- *ExtremeWare 7.2e Release Notes*
- *ExtremeWare 7.2e Command Reference Guide*

Documentation for Extreme Networks products is available on the World Wide Web at the following location:

http://www.extremenetworks.com/

Using ExtremeWare Publications Online

You can access ExtremeWare publications by downloading them from the Extreme Networks World Wide Web location or from your ExtremeWare product CD. Publications are provided in Adobe® Portable Document Format (PDF). Displaying or printing PDF files requires that your computer be equipped with Adobe® Reader® software, which is available free of charge from Adobe Systems Incorporated.

The following two ExtremeWare publications are available as PDF files that are designed to be used online together:

- *ExtremeWare 7.2e Installation and User Guide*
- *ExtremeWare 7.2e Command Reference Guide*

The user guide PDF file provides links that connect you directly to relevant command information in the command reference guide PDF file. This quick-referencing capability enables you to easily find detailed information in the command reference guide for any command mentioned in the user guide.

To ensure that the quick-referencing feature functions properly, follow these steps:

1. Download both the user guide PDF file and the command reference guide PDF file to the *same* destination directory on your computer.
2. You may open one or both PDF files and to enable cross-referenced linking between the user guide and command reference guide; however, it is recommended that for ease of use, you keep both files open concurrently on your computer desktop.

NOTE

If you activate a cross-referencing link from the ExtremeWare 7.2e Installation and User Guide PDF file to the command reference PDF file when the command reference PDF file is closed (that is, not currently open on your computer desktop), the system will close the user guide PDF file and open the command reference PDF file. To keep both PDF files open when you activate a cross-reference link, open both PDF files before using the link.

1 Summit 400-48 Switch Overview and Installation

This chapter describes the features and functionality of the Summit 400-48tes:

Summary of Features

Hardware

The Summit 400-48t supports the following ExtremeWare features:

- 48 copper ports 10/100/1000BASE-T
- 4 fiber SFP (mini-GBIC 1000BASE-SX, 1000BASE-LX, and 1000BASE-ZX)

The fiber ports share PHY with the first four copper port.

- 1 copper management port 10/100/1000BASE-T
- 1 console port, serial
- 2 (optional) modular 10 Gigabit uplink ports
- 2 stacking ports (10 Gigabit) reserved for future software features
- Supports redundant power support using the optional EPS 160 External Power Supply
- Redundant uplink support

Software

The software features of the Summit 400 include:

- Virtual local area networks (VLANs) including support for IEEE 802.1Q and IEEE 802.1p
- VLAN aggregation
- Spanning Tree Protocol (STP) (IEEE 802.1D)
- Quality of Service (QoS) including support for IEEE 802.1P, MAC QoS, and eight hardware queues
- Policy-Based Quality of Service (PB-QoS)
- Wire-speed Internet Protocol (IP) routing
- Extreme Standby Router Protocol (ESRP) - Aware support
- Ethernet Automated Protection Switching (EAPS) support
- IP Multinetting
- Jumbo frame support
- DHCP/BOOTP Relay
- Routing Information Protocol (RIP) version 1 and RIP version 2
- Open Shortest Path First (OSPF) routing protocol
- Wire-speed IP multicast routing support
- Diffserv support
- Access-policy support for routing protocols
- Access list support for packet filtering
- Access list support for rate-limiting
- IGMP snooping to control IP multicast traffic
- Protocol Independent Multicast-Sparse Mode (PIM-SM)
- Load sharing on multiple ports
- RADIUS client and per-command authentication support
- TACACS+ support
- Console command line interface (CLI) connection
- Telnet CLI connection
- SSH2 connection
- ExtremeWare Vista Web-based management interface
- Simple Network Management Protocol (SNMP) support

- Remote Monitoring (RMON)
- Traffic mirroring for ports by port number
- Network Login—Web
- Network Login—IEEE 802.1X

Summit 400-48t Switch Physical Features

The Summit 400-48t switch is a compact enclosure (see Figure 1) one rack unit in height (1.73 inches or 44.0 mm) that provides 48 autosensing 10/100/1000BASE-T ports using RJ-45 connectors. The switch also has four fiber ports that allow Gigabit Ethernet uplink connections through Extreme 1000BASE-SX, 1000BASE-LX, or 1000BASE-ZX SFP mini-GBICs using LC connectors. The four fiber ports and the first four of the 10/100/1000BASE-T ports are designed as shared, or *combination ports* for uplink redundancy. When sharing ports, only the fiber port or only the copper port can be active at the same time. For more information on cabling and configuring this feature, see "Uplink Redundancy" on page 27.

Summit 400-48t Switch Front View

Figure 1 shows the Summit 400-48t switch front view.

Figure 1: Summit 400-48t switch front view

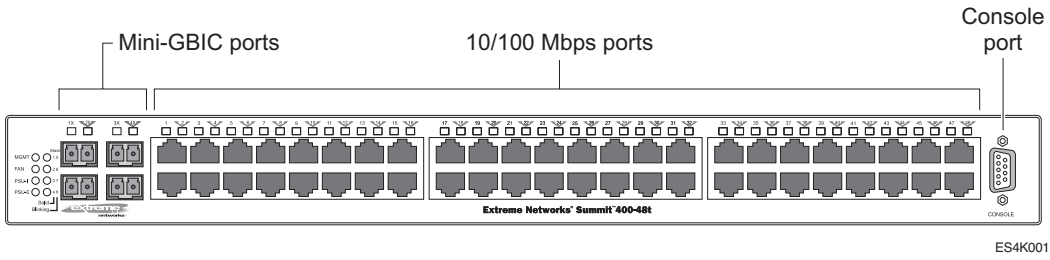

The front panel consists of:

LEDs—For a description of the LEDs and their behavior, see "Summit 400-48t Switch LEDs" on page 23.

Fiber uplink ports—For more information about these four ports, see "Mini-GBIC Type and Support" on page 24.

10/100/1000BASE-T ports—For more information about these 48 ports, see "Port Connections" on page 27.

Console Port—Use the console port (9-pin, "D" type connector) to attach a terminal and access the CLI through a serial connection. Use the console port to carry out local management.

Summit 400-48 Switch Rear View

Figure 2 shows the rear view of the Summit 400-48 switch.

Figure 2: Summit 400-48 switch rear view

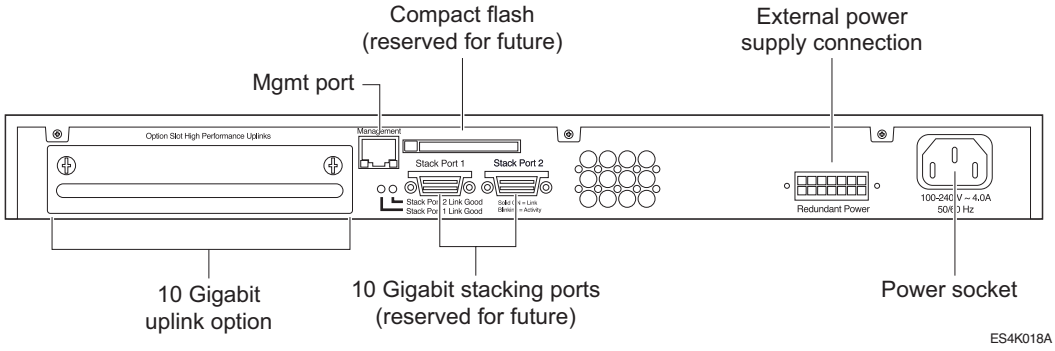

The rear panel consists of:

- An option slot for the dual 10 Gigabit uplinks

  To install this option, see “Installing Optional Features” on page 39.

- The management port

  The 10/100/1000BASE-T Ethernet management port communicates directly with the CPU of the switch, bypassing the switch. Connect an Ethernet cable directly from a laptop into the management port to view and locally manage the switch configurations.

  Do not assign an in-band IP address to the management port VLAN. The management port VLAN is an out-of-band VLAN, so if it is assigned an in-band IP address (an address where the source and destination are in the same subnet), the switch treats it as a normal VLAN and attempts to route traffic through it.

  Extreme Networks does not recommend that you use the management port to route traffic to any front panel port on the switch. The management port is designed only for switch management purposes.

  There are two LEDs for the management port, located in the bottom corners of the port. The LED on the bottom right turns solid green when a cable is inserted and the port detects a link. The LED on the bottom left blinks green when there is transmission activity on the link.

- A compact flash slot

  This slot is currently not supported but is reserved for future use.

- Two high-performance stacking ports

  These ports are currently not supported but are reserved for future software features.

- Vents for the internal power supply fan.

- The connector for the optional Extreme External Power Supply System.

  For further information about this feature, see “Installing Optional Features” on page 39.

- AC Power Socket

  The Summit 400-48 switch automatically adjusts to the supply voltage. The power supply operates from 100 VAC to 240 VAC.

NOTE

The Summit 400-48 switch certification, safety label, and serial number are located on the bottom of the switch.

Summit 400-48t Switch LEDs

The front panel displays five types of LEDs:

- Management

  The MGMT LED indicates the status of the switch.

- Fan

  The FAN LED indicates the status of the cooling fans.

- Power

  The Summit 400-48t comes with an internal power supply and can be connected to the Extreme External Power Supply tray. The status of the internal power supply is indicated by the PSU-I LED. The status of the external power supply is indicated by the PSU-E LED.

- 10/100/1000BASE-T port status

  Each of the 48 copper 10/100/1000BASE-T ports has an associated LED located above the port.

- Fiber port status

  Each of the four optical fiber ports has an associated LED located above the port.

Table 3 describes the behavior of the front panel LEDs on the Summit 400-48t switch.

Table 3: Summit 400-48t switch LED behavior

Unit Status LED (MGMT LED)		
	Color	**Indicates**
	Green, slow blinking	The Summit switch is operating normally.
	Green, fast blinking	The Summit switch POST is in progress.
	Green, solid	POST passed; ExtremeWare is booting.
	Amber, blinking	The Summit switch has failed its POST or an overheat condition is detected.
	Off	The Summit switch has no power.
Fan LED		
	Color	**Indicates**
	Green, solid	All fans are operating normally.
	Amber, blinking	One or more fans has failed. The switch continues to operate unless over-heating occurs.
	Off	The Summit switch has no power.

Table 3: Summit 400-48t switch LED behavior (Continued)

Power Supply LEDs		
PSU-I	**Color**	**Indicates**
	Green, solid	The internal power supply is operating normally.
	Amber, blinking	The internal power supply has failed. Replace the internal power supply as soon as possible.
	Off	The internal power supply has no power.
PSU-E	**Color**	**Indicates**
	Green, solid	The external power supply is operating normally.
	Amber, blinking	The external powersupply has failed. Replace the external power supply as soon as possible.
	Off	The external power supply is not connected.
Port Status LEDs (Ports 1–48)		
	Color	**Indicates**
	Green, solid	The link is present; port is enabled.
	Green blinking	The link is present and the port is transmitting or receiving packets.
	Off	The link is not present.
Fiber LEDs (Ports 1X—4X)		
	Color	**Indicates**
	Green, solid	Fiber link is selected; mini-GBIC is present and being used for the Gigabit Ethernet uplink.
	Green, blinking	The link is present and the port is transmitting or receiving packets.
	Off	1000BASE-T link is selected; the switch is using the RJ-45 port for the Gigabit Ethernet uplink.
Stack LEDs (Reserved for future features)		

Mini-GBIC Type and Support

The Summit 400-48t supports the SFP GBIC, also known as the mini-GBIC, in three types: the SX mini-GBIC, which conforms to the 1000BASE-SX standard, the LX mini-GBIC, which conforms to the 1000BASE-LX standard, and the ZX mini-GBIC, a long-haul mini-GBIC that conforms to the IEEE 802.3z standard. The system uses identifier bits to determine the media type of the mini-GBIC that is installed. The Summit 400-48tes support only the SFP mini-GBIC.

NOTE

Only mini-GBICs that have been certified by Extreme Networks (available from Extreme Networks) should be inserted into the mini-GBIC receptacles on the Summit 400-48t.

This section describes the mini-GBIC types and specifications.

Mini-GBIC Type and Specifications

Table 4 describes the mini-GBIC type and distances for the Summit 400-48t.

Table 4: Mini-GBIC types and distances

Standard	Media Type	Mhz•Km Rating	Maximum Distance (Meters)
1000BASE-SX (850 nm optical window)	50/125 µm multimode fiber	400	500
	50/125 µm multimode fiber	500	550
	62.5/125 µm multimode fiber	160	220
	62.5/125 µm multimode fiber	200	275
1000BASE-LX (1310 nm optical window)	50/125 µm multimode fiber	400	550
	50/125 µm multimode fiber	500	550
	62.5/125 µm multimode fiber	500	550
	10/125 µm single-mode fiber	—	5,000
1000BASE-ZX (1550 nm optical window)	10/125 µm single-mode fiber	—	50,000

SX Mini-GBIC Specifications

Table 5 describes the specifications for the SX mini-GBIC.

Table 5: SX mini-GBIC specifications

Parameter	Minimum	Typical	Maximum
Transceiver			
Optical output power	–9.5 dBm		–4 dBm
Center wavelength	830 nm	850 nm	860 nm
Receiver			
Optical input power sensitivity	–21 dBm		
Optical input power maximum			–4 dBm
Operating wavelength	830 nm		860 nm
General			
Total system budget			11.5 dB

Total optical system budget for the SX mini-GBIC is 11.5 dB. Extreme Networks recommends that 3 dB of the total budget be reserved for losses induced by cable splices, connectors, and operating margin. While 8.5 dB remains available for cable-induced attenuation, the 1000BASE-SX standard specifies supported distances of 275 meters over 62.5 micron multimode fiber and 550 meters over 50 micron multimode fiber. There is no minimum attenuation or minimum cable length restriction.

LX Mini-GBIC Specifications

Table 6 describes the specifications for the LX mini-GBIC.

Table 6: LX mini-GBIC specifications

Parameter	Minimum	Typical	Maximum
Transceiver			
Optical output power	–9.5 dBm		–3 dBm
Center wavelength	1275 nm	1310 nm	1355 nm
Receiver			
Optical input power sensitivity	–23 dBm		
Optical input power maximum			–3 dBm
Operating wavelength	1270 nm		1355 nm
General			
Total system budget			13.5 dB

Total optical system budget for the LX mini-GBIC is 13.5 dB. Measure cable plant losses with a 1310 nm light source and verify this to be within budget. When calculating the maximum distance attainable using optical cable with a specified loss per kilometer (for example 0.25 dB/km) Extreme Networks recommends that 3 dB of the total budget be reserved for losses induced by cable splices, connectors, and operating margin. Thus, 10.5 dB remains available for cable induced attenuation. There is no minimum attenuation or minimum cable length restriction.

ZX Mini-GBIC Specifications

Table 7 describes the specifications for the ZX mini-GBIC.

Table 7: ZX mini-GBIC specifications

Parameter	Minimum	Typical	Maximum
Transceiver			
Optical output power	–2 dBm	0 dBm	3 dBm
Center wavelength	1540 nm	1550 nm	1570 nm
Receiver			
Optical input power sensitivity	–23 dBm		
Optical input power maximum			–3 dBm
Operating wavelength	1540 nm	1550 nm	1570 nm

Long Range GBIC System Budgets

Measure cable plant losses with a 1550 nm light source and verify this to be within budget. When calculating the maximum distance attainable using optical cable with a specified loss per kilometer (for example 0.25 dB/km), Extreme Networks recommends that 3 dB of the total budget be reserved for losses induced by cable splices, connectors, and operating margin. Figure 3 shows the total optical system budget between long range GBICs in various end-to-end combinations (ZX, ZX Rev 03, LX70, and LX100).

NOTE

The ZX mini-GBIC is equivalent to the ZX Rev 03 GBIC.

Figure 3: Total optical system budgets for long range GBICs

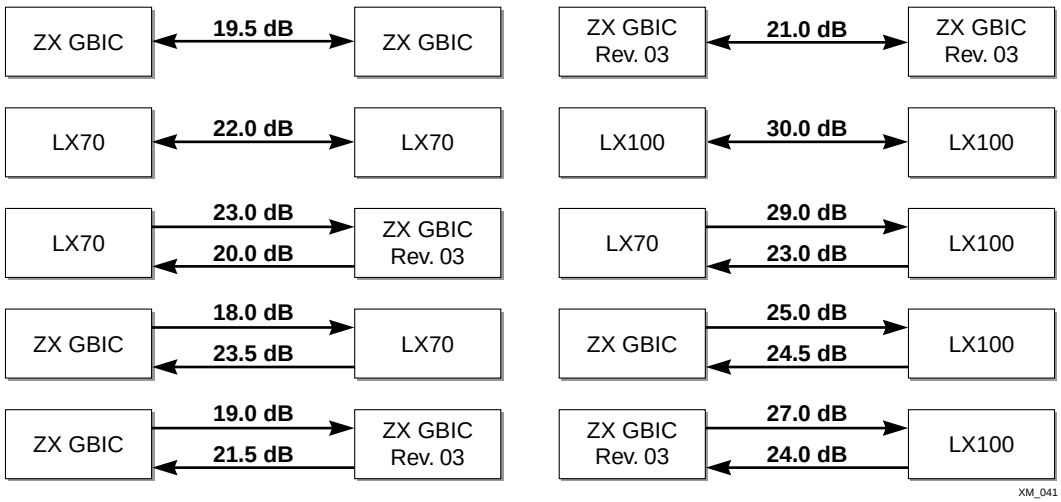

Table 8 lists the minimum attenuation requirements to prevent saturation of the receiver for each type of long range GBIC.

Table 8: Minimum attenuation requirements

		Receivers				
	GBIC Type	**LX70**	**LX100**	**ZX (prior to Rev 03)**	**ZX Rev 03**	**ZX mini**
Transceivers	**LX70**	9 dB	13 dB	7 dB	7 dB	9 dB
	LX100	8 dB	12 dB	6 dB	6 dB	8 dB
	ZX (prior to Rev 03)	2 dB	6 dB	0 dB	0 dB	2 dB
	ZX Rev 03	5 dB	9 dB	3 dB	3 dB	5 dB
	ZX mini	6 dB	10 dB	4 dB	4 dB	6 dB

Port Connections

The Summit 400-48t switch has 48 copper 10/100/1000BASE-T ports using RJ-45 connectors for communicating with end stations and other devices over 10/100/1000 Mbps Ethernet.

The switch provides full-duplex support for all ports. Full-duplex allows frames to be transmitted and received simultaneously and, in effect, doubles the bandwidth available on a link. All 10/100/1000 Mbps ports on the Summit 400-48t switch autonegotiate for half- or full-duplex operation.

Uplink Redundancy

The four fiber ports and the first four of the 10/100/1000BASE-T ports are designed as combination ports for uplink redundancy. When sharing ports, only the fiber port or only the copper port can be active at the same time. If copper port 1 goes down while transmitting packets, fiber port 1X activates and becomes the primary link. See Figure 4 for a diagram of these combination ports.

The switch determines whether the port is the primary or redundant port based upon the order in which the cables are inserted into the switch. When the switch senses that cables are in both the fiber and corresponding copper port, the switch enables the uplink redundancy feature. For example, if you insert mini-GBICs into ports 1X and 3X first, and then connect copper ports 1 and 3, the switch assigns ports 1 and 3 as redundant ports.

Figure 4: Redundancy cabling

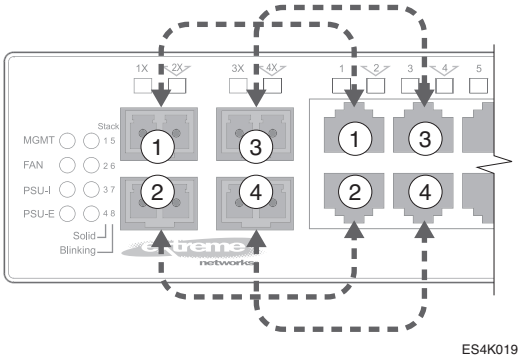

You can override the configuration and behavior of these ports through the CLI. Using the CLI, you can set a preference for either fiber or copper. You can also turn off port redundancy using the *force* option. If a combination port fails to link, determine whether the *force* option is in effect. For more information about using the CLI to set redundancy priority, see "Configuring Ports" on page 77.

The Summit 400-48 switch Gigabit Ethernet port failover from the fiber link to the copper link takes 4-5 seconds. The Summit 400-48t switch Gigabit Ethernet port failover from the copper link to the fiber link takes 2-3 seconds.

NOTE

To support automatic failover between the fiber and copper ports, you must use an Extreme mini-GBIC connector.

Software Overview

Virtual LANs (VLANs)

ExtremeWare has a VLAN feature that enables you to construct your broadcast domains without being restricted by physical connections. A VLAN is a group of location- and topology-independent devices that communicate as if they were on the same physical local area network (LAN).

Implementing VLANs on your network has the following three advantages:

- VLANs help to control broadcast traffic. If a device in VLAN *Marketing* transmits a broadcast frame, only VLAN *Marketing* devices receive the frame.
- VLANs provide extra security. Devices in VLAN *Marketing* can only communicate with devices on VLAN *Sales* using routing services.
- VLANs ease the change and movement of devices on networks.

For more information on VLANs, see Chapter 5.

Spanning Tree Protocol

The switch supports the IEEE 802.1D Spanning Tree Protocol (STP), which is a bridge-based mechanism for providing fault tolerance on networks. STP enables you to implement parallel paths for network traffic, and ensure that:

- Redundant paths are disabled when the main paths are operational.
- Redundant paths are enabled if the main traffic paths fail.

A single spanning tree can span multiple VLANs.

For more information on STP, see Chapter 11.

Quality of Service

ExtremeWare has Policy-Based Quality of Service (QoS) features that enable you to specify service levels for different traffic groups. By default, all traffic is assigned the *normal* QoS policy profile. If needed, you can create other QoS policies and apply them to different traffic types so that they have different guaranteed minimum bandwidth, maximum bandwidth, and priority. For more information on Quality of Service, see Chapter 7.

Unicast Routing

The switch can route IP traffic between the VLANs that are configured as virtual router interfaces. Both dynamic and static IP routes are maintained in the routing table. The following routing protocols are supported:

- RIP version 1
- RIP version 2
- OSPF version 2

For more information on IP unicast routing, see Chapter 12.

IP Multicast Routing

The switch can use IP multicasting to allow a single IP host to transmit a packet to a group of IP hosts. ExtremeWare supports multicast routes that are learned by way of the Protocol Independent Multicast (sparse mode). For more information on IP multicast routing, see Chapter 14.

Load Sharing

Load sharing allows you to increase bandwidth and resiliency by using a group of ports to carry traffic in parallel between systems. The load sharing algorithm allows the switch to use multiple ports as a single logical port. For example, VLANs see the load-sharing group as a single virtual port. The algorithm also guarantees packet sequencing between clients. For more information on load sharing, see Chapter 4.

ESRP-Aware Switches

Extreme switches that are not running ESRP, but are connected on a network that has other Extreme switches running ESRP are ESRP-*aware*. When ESRP-aware switches are attached to ESRP-enabled switches, the ESRP-aware switches reliably perform fail-over and fail-back scenarios in the prescribed recovery times. No configuration of this feature is necessary.

NOTE

If you disable EDP on the switch, the switch is no longer ESRP-aware.

If Extreme switches running ESRP are connected to layer 2 switches that are not manufactured by Extreme Networks (or Extreme switches that are not running ExtremeWare 4.0 or later), the fail-over times seen for traffic local to the segment may appear longer, depending on the application involved and the FDB timer used by the other vendor's layer 2 switch. As such, ESRP can be used with layer 2 switches from other vendors, but the recovery times vary.

The VLANs associated with the ports connecting an ESRP-aware switch to an ESRP-enabled switch must be configured using an 802.1Q tag on the connecting port, or, if only a single VLAN is involved, as untagged.

To display ESRP-aware information, use the following command:

```
show esrp-aware [vlan <vlan name>]
```

The display includes the group number, MAC address for the master of the group, and age of the information.

Software Licensing

Some Extreme Networks products have capabilities that are enabled by using a license key. Keys are typically unique to the switch, and are not transferable. Keys are stored in NVRAM and, once entered, persist through reboots, software upgrades, and reconfigurations. The following sections describe the features that are associated with license keys.

Router Licensing

Some switches support software licensing for different levels of router functionality. In the Summit 400-48t, routing protocol support is separated into two sets: Edge and Advanced Edge. Edge is a subset of Advanced Edge.

Edge Functionality

Edge functionality requires *no license key*. Extreme switches that ship with an Edge license, do not require a license key. Edge functionality includes all switching functions, and also includes all available layer 3 QoS, access list, and ESRP functions. L3 routing functions include support for:

- IP routing using RIP version 1 and/or RIP version 2
- IP routing between directly attached VLANs
- IP routing using static routes
- ESRP-aware

- Layer 3 QoS
- Access Lists, except rate limiting
- Network Login, both web-based and 802.1X

Advanced Edge Functionality

The Advanced Edge license enables support of additional routing protocols and functions, including:

- IP routing using OSPF
- IP multicast routing using PIM (Sparse Mode)
- EAPS-Edge

Product Support

The Summit 400 can support Advanced Edge functionality. However, the switch is enabled and shipped with an Edge license.

Verifying the Switch License

To verify the license, use the `show switch` command.

Obtaining an Advanced Edge License Voucher

You can order the desired functionality from the factory, using the appropriate model of the desired product. If you order licensing from the factory, the license arrives in a separate package from the switch. After the license key is installed, it should not be necessary to enter the information again. However, we recommend keeping the certificate for your records.

You can upgrade the licensing of an existing product by purchasing a voucher for the desired product and functionality. Please contact your supplier to purchase a voucher.

The voucher contains information and instructions on obtaining a license key for the switch using the Extreme Networks Support website at:

http://www.extremenetworks.com/support/techsupport.asp

or by phoning Extreme Networks Technical Support at:

- (800) 998-2408
- (408) 579-2826

Security Licensing

Certain additional ExtremeWare security features, such as the use of Secure Shell (SSH2) encryption, may be under United States export restriction control. Extreme Networks ships these security features in a disabled state. You can obtain information on enabling these features at no charge from Extreme Networks.

Obtaining a Security License

To obtain information on enabling features that require export restriction, access the Extreme Networks Support website at:

http://www.extremenetworks.com/go/security.htm

Fill out a contact form to indicate compliance or noncompliance with the export restrictions. If you are in compliance, you will be given information that will allow you to enable security features.

Security Features Under License Control

Summit 400-48t software supports the SSH2 protocol. SSH2 allows the encryption of Telnet session data between an SSH2 client and an Extreme Networks switch. The software also enables the switch to function as an SSH2 client, sending encrypted data to an SSH2 server on a remote system. This version of software also supports the Secure Copy Protocol (SCP). The encryption methods used are under U.S. export restriction control.

Software Factory Defaults

Table 9 shows factory defaults for global Summit 400-48t features.

Table 9: Summit 400-48t Global Factory Defaults

Item	Default Setting
Serial or Telnet user account	*admin* with no password and *user* with no password
Web network management	Enabled
Telnet	Enabled
SSH2	Disabled
SNMP	Enabled
SNMP read community string	*public*
SNMP write community string	*private*
RMON	Disabled
BOOTP	Enabled on the default VLAN (*default*)
QoS	All traffic is part of the default queue
QoS monitoring	Automatic roving
802.1p priority	Recognition enabled
Virtual LANs	Three VLANs predefined. VLAN named *default* contains all ports and belongs to the STPD named *s0.* VLAN *mgmt* exists only on switches that have an Ethernet management port, and contains only that port. The Ethernet management port is DTE only, and is not capable of switching or routing. VLAN *MacVLanDiscover* is used only when using the MAC VLAN feature.
802.1Q tagging	All packets are untagged on the default VLAN (*default).*
Spanning Tree Protocol	Disabled for the switch; enabled for each port in the STPD.
Forwarding database aging period	300 seconds (5 minutes)
IP Routing	Disabled
RIP	Disabled
OSPF	Disabled
IP multicast routing	Disabled
IGMP	Enabled
IGMP snooping	Enabled

Table 9: Summit 400-48t Global Factory Defaults (Continued)

Item	Default Setting
PIM-SM	Disabled
NTP	Disabled
DNS	Disabled
Port mirroring	Disabled

NOTE

For default settings of individual Summit 400-48t-features, see individual chapters in this guide.

Switch Installation

CAUTION

Use of controls or adjustments of performance or procedures other than those specified herein can result in hazardous radiation exposure.

Determining the Switch Location

The Summit 400-48t is suited for use in the office, where it can be free-standing or mounted in a standard 19-inch equipment rack. Alternately, the device can be rack-mounted in a wiring closet or equipment room. Two mounting brackets are supplied with the switch.

When deciding where to install the switch, ensure that:

- The switch is accessible and cables can be connected easily.
- Water or moisture cannot enter the case of the unit.
- Air-flow around the unit and through the vents in the side of the case is not restricted. You should provide a minimum of 1 inch (25 mm) clearance.
- No objects are placed on top of the unit.
- Units are not stacked more than four high if the switch is free-standing.

Following Safety Information

Before installing or removing any components of the switch, or before carrying out any maintenance procedures, read the safety information provided in this guide.

Installing the Switch

The Summit 400-48t can be mounted in a rack, or placed free-standing on a tabletop.

Rack Mounting

CAUTION

Do not use the rack mount kits to suspend the switch from under a table or desk, or to attach the switch to a wall.

To rack mount the Summit 400-48t:

1. Place the switch upright on a hard flat surface, with the front facing you.
2. Remove the existing screws from the sides of the case (retain the screws for Step 4).
3. Locate a mounting bracket over the mounting holes on one side of the unit.
4. Insert the screws and fully tighten with a suitable screwdriver, as shown in Figure 5.

Figure 5: Fitting the mounting bracket

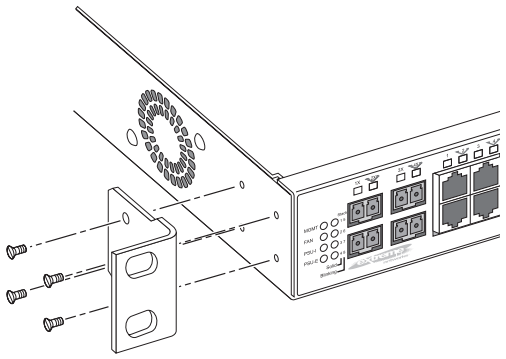

5. Repeat steps 2 through 4 for the other side of the switch.
6. Insert the switch into the 19-inch rack.
7. Secure the switch with suitable screws (not provided).
8. Connect the switch to the redundant power supply (if applicable). For further details of installing this option, see “Installing the External Power System” on page 42.
9. Connect cables.

Free-Standing

The Summit 400-48t is supplied with four self-adhesive rubber pads. Apply the pads to the underside of the device by sticking a pad in the marked area at each corner of the switch.

Desktop Mounting of Multiple Switches

You can physically place up to four Summit 400-48 switches on top of one another.

NOTE

This relates only to stacking the devices directly one on top of one another.

Apply the pads to the underside of the device by sticking a pad at each corner of the switch. Place the devices on top of one another, ensuring that the corners align.

Installing or Replacing a Mini-Gigabit Interface Connector (Mini-GBIC)

This section describes the safety precautions and preparation steps that you must perform before inserting and securing a mini-GBIC.

Safety Information

Before you install or replace a mini-GBIC, read the safety information in this section.

WARNING!

Mini-GBICs can emit invisible laser radiation. Avoid direct eye exposure to beam.

Mini-GBICs are a class 1 laser device. Use only devices approved by Extreme Networks. If a non-supported device is detected, a message is written to the syslog.

NOTE

Remove the LC fiber-optic connector from the mini-GBIC prior to removing the mini-GBIC from the switch.

Preparing to Install or Replace a Mini-GBIC

To ensure proper installation, complete the following tasks before inserting the mini-GBIC:

- Disable the port that is needed to install or replace the mini-GBIC.
- Inspect and clean the fiber tips, coupler, and connectors.
- Prepare and clean an external attenuator, if needed.
- Do not stretch the fiber.
- Make sure the bend radius of the fiber is not less than 2 inches.

In addition to the previously described tasks, Extreme Networks recommends the following when installing or replacing mini-GBICs on an active network:

- Use the same type of mini-GBIC at each end of the link.
- Connect one end of the link to the Tx port. Without an attenuator, measure the total loss from the Tx port to the other side of the link.

Once you complete all of the described tasks, you are ready to install or replace a mini-GBIC.

Removing and Inserting a Mini-GBIC

You can remove mini-GBICs from, or insert mini-GBICs into your Summit 400-48t without powering off the system. Figure 6 shows the two types of mini-GBIC modules.

Figure 6: Mini-GBIC modules

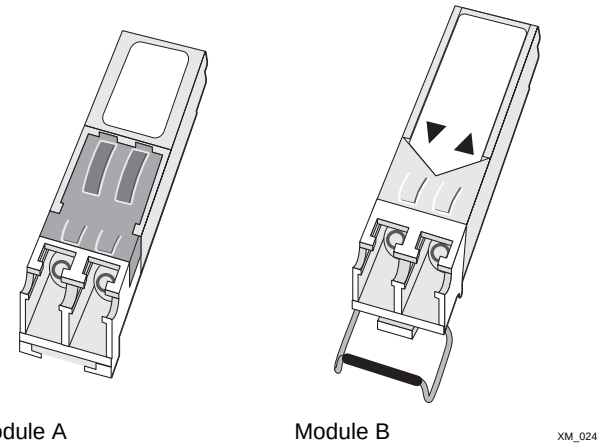

Mini-GBICs are a 3.3 V Class 1 laser device. Use only devices approved by Extreme Networks.

WARNING!

Mini-GBICs can emit invisible laser radiation. Avoid direct eye exposure to beam.

NOTE

Remove the LC fiber-optic connector from the mini-GBIC prior to removing the mini-GBIC from the switch.

Removing a Mini-GBIC

To remove a mini-GBIC similar to the one labeled “Module A” in Figure 6, gently press and hold the black plastic tab at the bottom of the connector to release the mini-GBIC, and pull the mini-GBIC out of the SFP receptacle on the switch.

To remove a mini-GBIC similar to the one labeled “Module B” in Figure 6, rotate the front handle down and pull the mini-GBIC out of the slot.

Inserting a Mini-GBIC

NOTE

Mini-GBICs can be installed in the SFP mini-GBIC receptacles for ports 1X—4X on the Summit 400-48tes.

To insert a mini-GBIC connector:

1. Holding the mini-GBIC by its sides, insert the mini-GBIC into the SFP receptacle on the switch.
2. Push the mini-GBIC into the SFP receptacle until you hear an audible click, indicating the mini-GBIC is securely seated in the SFP receptacle. If the mini-GBIC has a handle, push up on the handle to secure the mini-GBIC.

Connecting Equipment to the Console Port

Connection to the console port is used for direct local management. The switch console port settings are set as follows:

- **Baud rate**—9600
- **Data bits**—8
- **Stop bit**—1
- **Parity**—None
- **Flow control**—None

NOTE

If you set the switch console port flow control to XON/XOFF rather than None, you will be unable to access the switch. Do not set the switch console port flow control to XON/XOFF.

The terminal connected to the console port on the switch must be configured with the same settings. This procedure is described in the documentation supplied with the terminal.

Appropriate cables are available from your local supplier. To make your own cables, pinouts for a DB-9 male console connector are described in Table 10.

Table 10: Console Connector Pinouts

Function	Pin Number	Direction
DCD (data carrier detect)	1	In
RXD (receive data)	2	In
TXD (transmit data)	3	Out
DTR (data terminal ready)	4	Out
GND (ground)	5	—
DSR (data set ready)	6	In
RTS (request to send)	7	Out
CTS (clear to send	8	In
Not Connected	9	

Figure 7 shows the pin-outs for a 9-pin to RS-232 25-pin null-modem cable.

Figure 7: Null-modem cable pin-outs

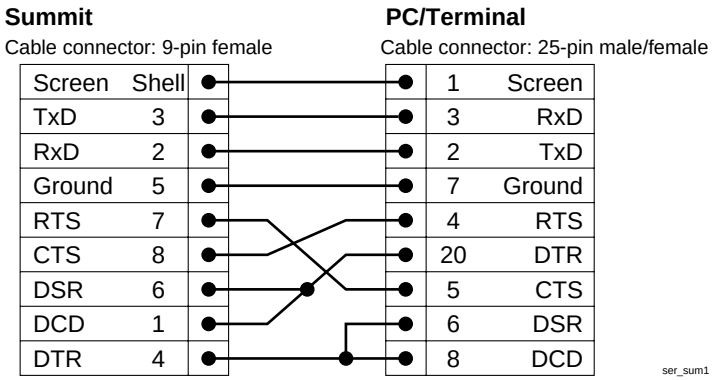

Figure 8 shows the pin-outs for a 9-pin to 9-pin PC-AT null-modem serial cable.

Figure 8: PC-AT serial null-modem cable pin-outs

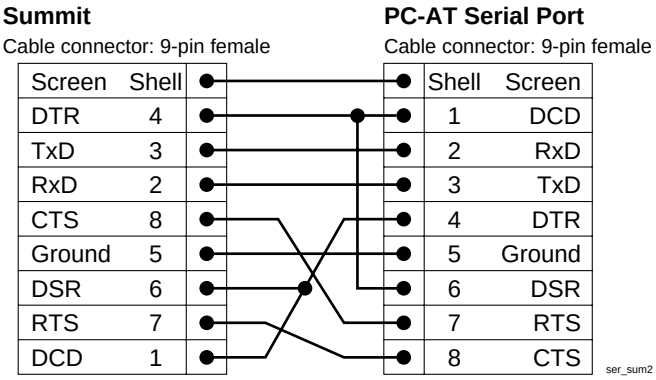

Powering On the Switch

To turn on power to the switch, connect the AC power cable to the switch and then to the wall outlet.

Checking the Installation

After turning on power to the Summit 400-48t, the device performs a Power On Self-Test (POST).

During the POST, all ports are temporarily disabled, the port LED is off, and the MGMT LED flashes fast. The MGMT LED flashes until the switch successfully passes the POST.

If the switch passes the POST, the MGMT LED is blinking slowly (once per second). If the switch fails the POST, the MGMT LED is amber. For more information on the LEDs, see “Summit 400-48 Switch Rear View” on page 22.

Logging In for the First Time

After the Summit 400-48t completes the POST, it is operational. Once operational, you can log in to the switch and configure an IP address for the default VLAN (named *default*).

To configure the IP settings manually, follow these steps:

1. Connect a terminal or workstation running terminal-emulation software to the console port.
2. At your terminal, press [Return] one or more times until you see the login prompt.
3. At the login prompt, enter the default user name *admin* to log on with administrator privileges.

   For example:

   ```
   login: admin
   ```

   Administrator capabilities allow you to access all switch functions.

   For more information on switch security, see "Network Login" on page 146.
4. At the password prompt, press [Return].

   The default name, *admin*, has no password assigned. When you have successfully logged on to the switch, the command-line prompt displays the name of the switch (for example, *Summit 400-48t*) in its prompt.
5. Assign an IP address and subnetwork mask for VLAN *default* by typing

   ```
   config vlan default ipaddress 123.45.67.8 255.255.255.0
   ```

   Your changes take effect immediately.
6. Save your configuration changes so that they will be in effect after the next switch reboot, by using the following command:

   ```
   save configuration {primary | secondary}
   ```

   For more information on saving configuration changes, see "Saving Configuration Changes" on page 315.
7. When you are finished using the facility, logout of the switch by typing

   ```
   logout
   ```

   After two incorrect login attempts, the Summit 400-48t locks you out of the login facility. You must wait a few minutes before attempting to log in again.

Installing Optional Features

Extreme Networks offers two hardware products that extend the capabilities of the Summit 400-48t. The Summit XEN Card is an additional card that adds one or two 10 Gigabit uplink modules through the back of the Summit 400-48t. The External Power Supply System allows you to attach an external power supply for backup to the internal power supply in the Summit 400-48t. Both of these products are additional offerings and available from your sales representative.

Before installing any optional features, be sure to check the Installation Notes provided with the feature to determine the latest installation process or limitations.

Installing the Summit XEN Card

The Summit 400-48t allows you to add up to two 10 Gigabit uplink modules to increase the bandwidth of the switch. The Summit XEN Card supports either of these Extreme XENPAK optical transceivers:

- SR XENPAK for the 850 nm range
- LR XENPAK for the 1310 nm range
- ER XENPAK for the 1550 nm range

CAUTION

The Summit XEN Card cannot be hot-swapped. Before installing the Summit XEN Card into the Summit 400-48t, you must turn off the switch. Use only XENPAK modules approved by Extreme Networks.

To install the Summit XEN Card:

1. Disconnect the AC power from the Summit 400.
2. Use a standard screwdriver to remove the blank plate to expose the opening for the card.
3. Install the XENPAK optical transceiver modules onto the card. For more detail on this step, see "Installing XENPAK Optical Transceiver Modules".
4. Place the Summit XEN Card into the drawer.
5. Carefully close the drawer to engage the card.

Installing XENPAK Optical Transceiver Modules

This section describes installing and removing the XENPAK module, a 10 Gbps optical transceiver. Both the LR XENPAK and the ER XENPAK appear and install the same. An example of an XENPAK module is shown in Figure 9.

Figure 9: XENPAK Modules

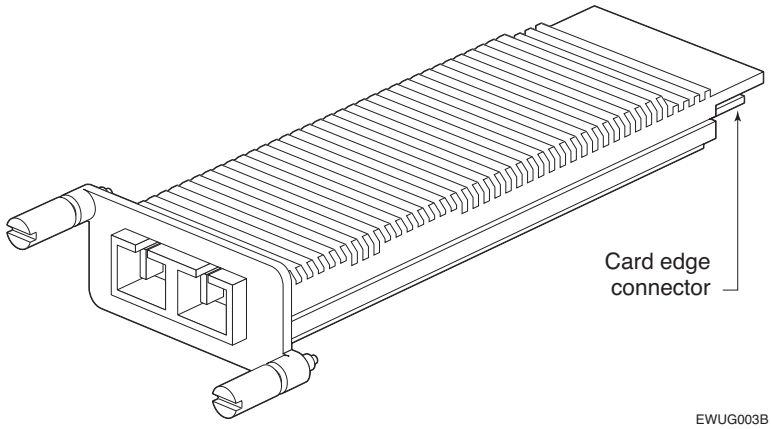

The XENPAK module is a Class 1 Laser device that operates at 5 V. Use only Extreme-approved devices on all Extreme switches.

CAUTION

The XENPAK module can emit invisible laser radiation. Avoid direct eye exposure to beam.

WARNING!

To prevent ESD damage to the Product Name, always use an ESD-preventive wrist strap when installing or removing the module. Handle the module by its sides only. Never touch the card-edge connectors at the insertion end of the module.

To install XENPAK modules:

1. Remove the XENPAK module from its antistatic container.
2. Remove the dust covers from the module connectors. If your module has a protective pad covering the card-edge connector, remove it.
3. Store the antistatic container, dust covers, and card-edge connector protective pad in a clean location in case you need to uninstall the module.
4. Hold the module by its sides and insert it into one of the two module slots on the Summit XEN card.
5. Slide the module as far back into the slot as possible, until you hear it click, indicating that it is firmly attached.
6. Secure the module to the card by turning the two captive screws clockwise until they are hand-tight.
7. Place the Summit XEN Card into the supplied drawer and carefully slide the drawer into the switch housing until the card seats and the drawer is flush with the remainder of the back panel.
8. Hand tighten the screws clockwise on the faceplate to keep the Summit XEN Card in place.

NOTE

To ensure that your module is undamaged upon installation, you can correlate factory test data with your installation site test data by consulting the average power reference values shown on the XENPAK module test data sheet (Part No. 121074-00) enclosed with your module.

To remove an XENPAK module:

1. Turn the two captive screws counter-clockwise until they are completely free from the Summit XEN. (The captive screws remain attached to the XENPAK module.)

   WARNING!

   Remove the SC fiber-optic connector from the XENPAK module before removing the module from the Summit XEN card.

2. Gripping both captive screws, pull the XENPAK module out of the card.
3. Place the dust covers back into the XENPAK module connectors.
4. Place the XENPAK module immediately into an antistatic container to protect it from ESD damage and dust.

Installing the External Power System

The Extreme External Power System allows you to add a redundant power supply to the Summit 400 in case of a power supply failure. It consists of a tray (EPS-T) that holds one or two EPS-160 power supplies, that provide one-to-one coverage for each External Power System that you attach.

To install the EPS-160:

1. Rack mount or place on a desktop near the External Power System. Use the procedure in "Rack Mounting" on page 34 help in installing the unit in a rack.
2. Connect the EPS-160 power supply to the back of the External Power System using the supplied cables.

   NOTE

   The cable length is 1 meter.

   The cable connector has a tab that fits into the chassis to ensure correct alignment of the connector. See Figure 10 to locate the connectors.

Figure 10: Redundant Power Connection

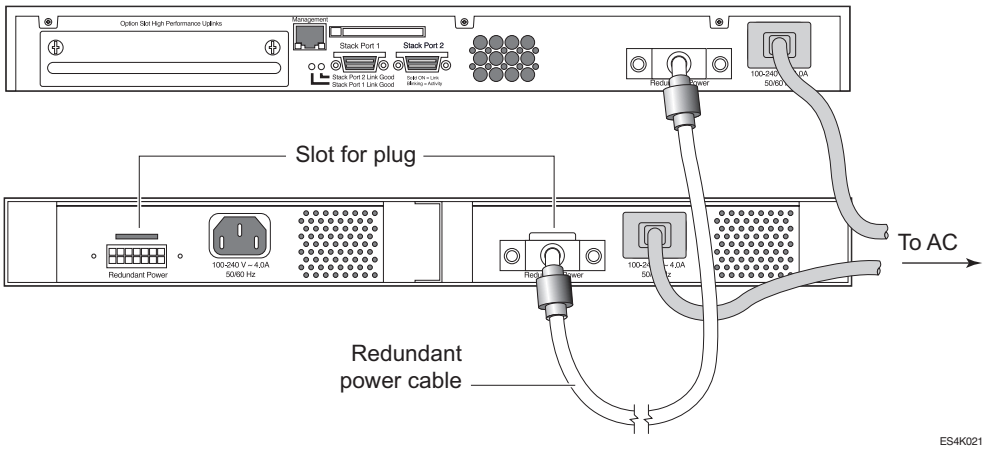

3. Connect the power to the External Power Supply System. The PSU-E LED on the front of the External Power System should be solid green to indicate that is ready. See Table 3 on page 23, for a full description of the LED indicators.

2 Managing the Switch

This chapter covers the following topics:

Overview

Using ExtremeWare, you can manage the switch using the following methods:

- Access the CLI by connecting a terminal (or workstation with terminal-emulation software) to the console port.
- Access the switch remotely using TCP/IP through one of the switch ports or through the dedicated 10/100/1000 unshielded twisted pair (UTP) Ethernet management port (on switches that are so equipped). Remote access includes:
  - Telnet using the CLI interface.
  - SSH2 using the CLI interface.
  - ExtremeWare Vista web access using a standard web browser.
  - SNMP access using EPICenter or another SNMP manager.
- Download software updates and upgrades. For more information, see Appendix B, Software Upgrade and Boot Options.

The switch supports up to the following number of concurrent user sessions:

- One console session
- Eight Telnet sessions
- Eight SSH2 sessions
- One web session

Using the Console Interface

The CLI built into the switch is accessible by way of the 9-pin, RS-232 port labeled *console*, located on the back of the switch. For more information on the console port pinouts, see Table 10 on page 37.

After the connection has been established, you will see the switch prompt and you can log in.

Using the 10/100/1000 Ethernet Management Port

The Summit 400 provides a dedicated 10/100/1000 Ethernet management port. This port provides dedicated remote access to the switch using TCP/IP. It supports the following management methods:

- Telnet using the CLI interface
- ExtremeWare Vista web access using a standard web browser
- SNMP access using EPICenter or another SNMP manager

The management port is a DTE port, and is not capable of supporting switching or routing functions. The TCP/IP configuration for the management port is done using the same syntax as used for VLAN configuration. The VLAN *mgmt* comes pre configured with only the 10/100/1000 UTP management port as a member.

You can configure the IP address, subnet mask, and default router for the VLAN *mgmt*, using the following commands:

```
configure vlan <vlan name> ipaddress <ipaddress> {<netmask> | <mask length>}
configure iproute add default <gateway> {<metric>}
```

Using Telnet

Any workstation with a Telnet facility should be able to communicate with the switch over a TCP/IP network using VT-100 terminal emulation.

Up to eight active Telnet sessions can access the switch concurrently. If `idletimeouts` are enabled, the Telnet connection will time out after 20 minutes of inactivity. If a connection to a Telnet session is lost inadvertently, the switch terminates the session within two hours.

Before you can start a Telnet session, you must set up the IP parameters described in “Configuring Switch IP Parameters” later in this chapter. Telnet is enabled by default.

NOTE

Maximize the Telnet screen so that automatically updating screens display correctly.

To open the Telnet session, you must specify the IP address of the device that you want to manage. Check the user manual supplied with the Telnet facility if you are unsure of how to do this.

After the connection is established, you will see the switch prompt and you may log in.

Connecting to Another Host Using Telnet

You can Telnet from the current CLI session to another host using the following command:

```
telnet [<ipaddress> | <hostname>] {<port_number>}
```

If the TCP port number is not specified, the Telnet session defaults to port 23. Only VT100 emulation is supported.

Configuring Switch IP Parameters

To manage the switch by way of a Telnet connection or by using an SNMP Network Manager, you must first configure the switch IP parameters.

Using a BOOTP Server

If you are using IP and you have a Bootstrap Protocol (BOOTP) server set up correctly on your network, you must provide the following information to the BOOTP server:

- Switch Media Access Control (MAC) address, found on the rear label of the switch
- IP address
- Subnet address mask (optional)

After this is done, the IP address and subnet mask for the switch will be downloaded automatically. You can then start managing the switch using this addressing information without further configuration.

You can enable BOOTP on a per-VLAN basis by using the following command:

```
enable bootp vlan [<vlan name> | all]
```

By default, BOOTP is enabled on the *default* VLAN.

If you configure the switch to use BOOTP, the switch IP address is not retained through a power cycle, even if the configuration has been saved. To retain the IP address through a power cycle, you must configure the IP address of the VLAN using the command-line interface, Telnet, or web interface.

All VLANs within a switch that are configured to use BOOTP to get their IP address use the same MAC address. Therefore, if you are using BOOTP relay through a router, the BOOTP server relays packets based on the gateway portion of the BOOTP packet.

NOTE

For more information on DHCP/BOOTP relay, see Chapter 12.

Manually Configuring the IP Settings

If you are using IP without a BOOTP server, you must enter the IP parameters for the switch in order for the SNMP Network Manager, Telnet software, or web interface to communicate with the device. To assign IP parameters to the switch, you must perform the following tasks:

- Log in to the switch with administrator privileges using the console interface.
- Assign an IP address and subnet mask to a VLAN.

  The switch comes configured with a default VLAN named *default*. To use Telnet or an SNMP Network Manager, you must have at least one VLAN on the switch, and it must be assigned an IP address and subnet mask. IP addresses are always assigned to each VLAN. The switch can be assigned multiple IP addresses.

NOTE

For information on creating and configuring VLANs, see Chapter 5.

To manually configure the IP settings, follow these steps:

1. Connect a terminal or workstation running terminal-emulation software to the console port, as detailed in "Using the Console Interface" on page 44.
2. At your terminal, press [Return] one or more times until you see the login prompt.
3. At the login prompt, enter your user name and password. Note that they are both case-sensitive. Ensure that you have entered a user name and password with administrator privileges.
   - If you are logging in for the first time, use the default user name *admin* to log in with administrator privileges. For example:

     ```
     login: admin
     ```

     Administrator capabilities enable you to access all switch functions. The default user names have no passwords assigned.
   - If you have been assigned a user name and password with administrator privileges, enter them at the login prompt.
4. At the password prompt, enter the password and press [Return].

   When you have successfully logged in to the switch, the command-line prompt displays the name of the switch in its prompt.
5. Assign an IP address and subnetwork mask for the default VLAN by using the following command:

   ```
   configure vlan <vlan name> ipaddress <ipaddress> {<netmask> | <mask length>}
   ```

   For example:

   ```
   configure vlan default ipaddress 123.45.67.8 255.255.255.0
   ```

   Your changes take effect immediately.

   NOTE

   As a general rule, when configuring any IP addresses for the switch, you can express a subnet mask by using dotted decimal notation, or by using classless inter-domain routing notation (CIDR). CIDR uses a forward slash plus the number of bits in the subnet mask. Using CIDR notation, the command identical to the one above would be:

   ```
   configure vlan default ipaddress 123.45.67.8 / 24
   ```

6 Configure the default route for the switch using the following command:

```
configure iproute add default <gateway> {<metric>}
```

For example:

```
configure iproute add default 123.45.67.1
```

7 Save your configuration changes so that they will be in effect after the next switch reboot, by using the following command:

```
save configuration {primary | secondary}
```

8 When you are finished using the facility, log out of the switch by typing:

`logout` or `quit`

Disconnecting a Telnet Session

An administrator-level account can disconnect a Telnet management session. If this happens, the user logged in by way of the Telnet connection is notified that the session has been terminated.

To terminate a Telnet session, follow these steps:

1 Log in to the switch with administrator privileges.

2 Determine the session number of the session you want to terminate by using the following command:

```
show session
```

3 Terminate the session by using the following command:

```
clear session <number>
```

Controlling Telnet Access

By default, Telnet services are enabled on the switch. Telnet access can be restricted by the use of an access profile. An access profile permits or denies a named list of IP addresses and subnet masks. To configure Telnet to use an access profile, use the following command:

```
enable telnet {access-profile [<access_profile> | none]} {port <tcp_port_number>}
```

Use the none option to remove a previously configured access profile.

To display the status of Telnet, use the following command:

```
show management
```

You can choose to disable Telnet by using the following command:

```
disable telnet
```

To re-enable Telnet on the switch, at the console port use the following:

```
enable telnet
```

You must be logged in as an administrator to enable or disable Telnet.

NOTE

For more information on Access Profiles, see Chapter 9.

Using Secure Shell 2 (SSH2)

Secure Shell 2 (SSH2) is a feature of ExtremeWare that allows you to encrypt Telnet session data between a network administrator using SSH2 client software and the switch, or to send encrypted data from the switch to an SSH2 client on a remote system. Image and configuration files may also be transferred to the switch using the Secure Copy Protocol 2 (SCP2). The ExtremeWare CLI provides a command that enable the switch to function as an SSH2 client, sending commands to a remote system via an SSH2 session. It also provides commands to copy image and configuration files to the switch using the SCP2.

For detailed information about SSH2 and SCP2, see Chapter 9, “Security”.

Using SNMP

Any Network Manager running the Simple Network Management Protocol (SNMP) can manage the switch, provided the Management Information Base (MIB) is installed correctly on the management station. Each Network Manager provides its own user interface to the management facilities.

The following sections describe how to get started if you want to use an SNMP manager. It assumes you are already familiar with SNMP management. If not, refer to the following publication:

> *The Simple Book*
> by Marshall T. Rose
> ISBN 0-13-8121611-9
> Published by Prentice Hall.

Enabling and Disabling SNMPv1/v2c and SNMPv3

ExtremeWare versions since 7.1.0 can concurrently support SNMPv1/v2c and SNMPv3. The default for the switch is to have both types of SNMP enabled. Network managers can access the device with either SNMPv1/v2c methods or SNMPv3. To enable concurrent support, use the following command:

```
enable snmp access
```

To prevent any type of SNMP access, use the following command:

```
disable snmp access
```

To prevent access using SNMPv1/v2c methods and allow access using SNMPv3 methods only, use the following commands:

```
enable snmp access
disable snmp access {snmp-v1v2c}
```

There is no way to configure the switch to allow SNMPv1/v2c access and prevent SNMPv3 access.

Most of the commands that support SNMPv1/v2c use the keyword `snmp`, most of the commands that support SNMPv3 use the keyword `snmpv3`.

Accessing Switch Agents

To have access to the SNMP agent residing in the switch, at least one VLAN must have an IP address assigned to it.

By default, SNMP access and SNMPv1/v2c traps are enabled. SNMP access and SNMP traps can be disabled and enabled independently—you can disable SNMP access but still allow SNMP traps to be sent, or vice versa.

Supported MIBs

In addition to private MIBs, the switch supports the standard MIBs listed in Appendix A.

NOTE

The SNMP ifAdminStatus MIB value is not saved after a reboot. Ports set to down in the SNMP ifAdminStatus MIB come back after rebooting. However, if you save the configuration using the CLI or SNMP after changing the port status to down in the ifAdminStatus MIB, then the change is saved after a reboot.

Configuring SNMPv1/v2c Settings

The following SNMPv1/v2c parameters can be configured on the switch:

- **Authorized trap receivers**—An authorized trap receiver can be one or more network management stations on your network. The switch sends SNMPv1/v2c traps to all trap receivers. You can have a maximum of 16 trap receivers configured for each switch, and you can specify a community string and UDP port for individually for each trap receiver. All community strings must also be added to the switch using the `configure snmp add community` command.

  To configure a trap receiver on a switch, use the following command:

  ```
  configure snmp add trapreceiver <ip address> {port <number>} community {hex}
  <community string> {from <source ip address>} {mode [enhanced | standard]}
  trap-group {auth-traps{,}}  {extreme-traps{,}} {link-up-down-traps{,}}
  {ospf-traps{,} {ping-traceroute-traps{,}} {rmon-traps{,}} {security-traps{,}}
  {smart-traps{,}} {stp-traps{,}} {system-traps{,}} {vrrp-traps{,}}
  ```

  See the Command Reference for a listing of the available traps.

  You can delete a trap receiver using the `configure snmp delete trapreceiver` command.

  Entries in the trap receiver list can also be created, modified, and deleted using the RMON2 trapDestTable MIB variable, as described in RFC 2021.

- **SNMP read access**—The ability to read SNMP information can be restricted through the use of an access profile. An access profile permits or denies a named list of IP addresses and subnet masks.

  To configure SNMPv1/v2c read access to use an access profile, use the following command:

  ```
  configure snmp access-profile readonly [<access-profile> | none]
  ```

  Use the none option to remove a previously configured access profile.

- **SNMP read/write access**—The ability to read and write SNMP information can be restricted through the use of an access profile. An access profile permits or denies a named list of IP addresses and subnet masks.

  To configure SNMPv1/v2c read/write access to use an access profile, use the following command:

  ```
  configure snmp access-profile readwrite [<access-profile> | none]
  ```

Use the none option to remove a previously configured access-profile.

- **Community strings**—The community strings allow a simple method of authentication between the switch and the remote Network Manager. There are two types of community strings on the switch. Read community strings provide read-only access to the switch. The default read-only community string is *public*. Read-write community strings provide read and write access to the switch. The default read-write community string is *private*.
- **System contact** (optional)—The system contact is a text field that enables you to enter the name of the person(s) responsible for managing the switch.
- **System name**—The system name is the name that you have assigned to this switch. The default name is the model name of the switch (for example, Summit1 switch).
- **System location** (optional)—Using the system location field, you can enter an optional location for this switch.
- **Enabling/disabling link up and link down traps** (optional)—By default, link up and link down traps (also called port-up-down traps) are enabled on the switch for all ports. SNMPv1 traps for link up and link down are not supported; ExtremeWare uses SNMPv2 traps.

  You can disable or re-enable the sending of these traps on a per port basis, by using the following commands:

  ```
  disable snmp traps port-up-down ports [all | mgmt | <portlist>]
  enable snmp traps {port-up-down ports [all | mgmt | <portlist>]}
  ```

  The mgmt option will only appear on platforms having a management port.

Displaying SNMP Settings

To display the SNMP settings configured on the switch, use the following command:

```
show management
```

This command displays the following information:

- Enable / disable state for Telnet, SSH2, SNMP, and web access, along with access profile information
- SNMP community strings
- Authorized SNMP station list
- SNMP MAC-security traps
- Link up / link down traps enabled on ports
- SNMP trap receiver list
- SNMP trap groups
- RMON polling configuration
- Login statistics
- Enable / disable status of link up and link down traps
- Enable / disable status of MAC-security traps

SNMP Trap Groups

SNMP trap groups allow you to specify which SNMP traps to send to a particular trap receiver. This functionality was made possible by the underlying support for SNMPv3. Essentially, a number of predefined filters are associated with a trap receiver, so that only those traps are sent. If you have

already been using SNMPv1/v2c trap receivers, trap groups are very easy to incorporate into your network. You cannot define your own trap groups. If you need to define more selectively which notifications to receive, you will need to use the notification filter capabilities available in SNMPv3.

To configure trap groups, use the following command:

```
configure snmp add trapreceiver <ip address> {port <number>} community {hex}
<community string> {from <source ip address>} {mode [enhanced | standard]} trap-group
{auth-traps{,}}  {extreme-traps{,}} {link-up-down-traps{,}} {ospf-traps{,}
{ping-traceroute-traps{,}} {rmon-traps{,}} {security-traps{,}} {smart-traps{,}}
{stp-traps{,}} {system-traps{,}} {vrrp-traps{,}}
```

For example, to send system and link up/link down traps to the receiver at 10.20.30.44 port 9347 with the community string *private*, use the following command:

```
configure snmp add trapreceiver 10.20.30.44 port 9347 community private trap-group
link-up-down-traps , system-traps
```

Table 11 lists the currently defined SNMP trap groups. From time to time, new trap groups may be added to this command.

Table 11: SNMP Trap Groups

Trap Group	Notifications	MIB Subtree
stp-traps	newRoot topologyChange	dot1dBridge, 1.3.6.1.2.1.17
ospf-traps	ospfIfStateChange ospfVirtIfStateChange ospfNbrStateChange ospfVirtNbrStateChange ospfIfConfigError ospfVirtIfConfigError ospfIfAuthFailure ospfVirtIfAuthFailure ospfIfRxBadPacket ospfVirtIfRxBadPacket ospfTxRetransmit ospfVirtIfTxRetransmit ospfOriginateLsa ospfMaxAgeLsa ospfLsdbOverflow ospfLsdbApproachingOverflow	ospfTraps, 1.3.6.1.2.1.14.16.2
ping-traceroute-traps	pingTestFailed pingTestCompleted	pingNotifications, 1.3.6.1.2.1.80.0
	tracerouteTestFailed tracerouteTestCompleted	traceRouteNotifications, 1.3.6.1.2.1.81.0
vrrp-traps	vrrpTrapNewMaster vrrpTrapAuthFailure	vrrpNotifications, 1.3.6.1.2.1.68.0

Table 11: SNMP Trap Groups (Continued)

Trap Group	Notifications	MIB Subtree
system-traps	extremeOverheat	1.3.6.1.4.1.1916.0.6
	extremeFanFailed	1.3.6.1.4.1.1916.0.7
	extremeFanOK	1.3.6.1.4.1.1916.0.8
	extremePowerSupplyFail	1.3.6.1.4.1.1916.0.10
	extremePowerSupplyGood	1.3.6.1.4.1.1916.0.11
	extremeModuleStateChange	1.3.6.1.4.1.1916.0.15
	extremeHealthCheckFailed	1.3.6.1.4.1.1916.4.1.0.1
	extremeCpuUtilizationRisingTrap	1.3.6.1.4.1.1916.4.1.0.2
	extremeCpuUtilizationFallingTrap	1.3.6.1.4.1.1916.4.1.0.3
	coldStart	1.3.6.1.6.3.1.1.5.1
	warmStart	1.3.6.1.6.3.1.1.5.2
extreme-traps	extremeEsrpStateChange	1.3.6.1.4.1.1916.0.17
	extremeEdpNeighborAdded	1.3.6.1.4.1.1916.0.20
	extremeEdpNeighborRemoved	1.3.6.1.4.1.1916.0.21
	extremeSlbUnitAdded	1.3.6.1.4.1.1916.0.18
	extremeSlbUnitRemoved	1.3.6.1.4.1.1916.0.19
smart-traps	extremeSmartTrap	1.3.6.1.4.1.1916.0.14
auth-traps	AuthenticationFailure	1.3.6.1.6.3.1.1.5.5
	extremeInvalidLoginAttempt	1.3.6.1.4.1.1916.0.9
link-up-down-traps	linkDown	1.3.6.1.6.3.1.1.5.3
	linkUp	1.3.6.1.6.3.1.1.5.4
rmon-traps	risingAlarm fallingAlarm	rmon-traps, 1.3.6.1.2.1.16.0
security-traps	extremeMacLimitExceeded	1.3.6.1.4.1.1916.4.3.0.1
	extremeUnauthorizedPortForMacDetected	1.3.6.1.4.1.1916.4.3.0.2
	extremeMacDetectedOnLockedPort	1.3.6.1.4.1.1916.4.3.0.3
	extremeNetloginUserLogin	1.3.6.1.4.1.1916.4.3.0.4
	extremeNetloginUserLogout	1.3.6.1.4.1.1916.4.3.0.5
	extremeNetloginAuthFailure	1.3.6.1.4.1.1916.4.3.0.6

SNMPv3

Beginning in ExtremeWare version 7.1.0, support was added for SNMPv3. SNMPv3 is an enhanced standard for SNMP that improves the security and privacy of SNMP access to managed devices, and provides sophisticated control of access to the device MIB. The prior standard versions of SNMP, SNMPv1 and SNMPv2c provided no privacy and little (or no) security.

The following six RFCs provide the foundation for Extreme Networks implementation of SNMPv3:

- RFC 3410, *Introduction to version 3 of the Internet-standard Network Management Framework*, provides an overview of SNMPv3.
- RFC 3411, *An Architecture for Describing SNMP Management Frameworks*, talks about SNMP architecture, especially the architecture for security and administration.
- RFC 3412, *Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)*, talks about the message processing models and dispatching that can be a part of an SNMP engine.
- RFC 3413, *SNMPv3 Applications*, talks about the different types of applications that can be associated with an SNMPv3 engine.
- RFC 3414, *The User-Based Security Model for Version 3 of the Simple Network Management Protocol (SNMPv3)*, describes the User-Based Security Model (USM).

- RFC 3415, *View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)*, talks about VACM as a way to access the MIB.

SNMPv3 Overview

The SNMPv3 standards for network management were primarily driven the need for greater security and access control. The new standards use a modular design and model management information by cleanly defining a message processing subsystem, a security subsystem, and an access control subsystem.

The message processing (MP) subsystem helps identify the MP model to be used when processing a received Protocol Data Unit (PDU), the packets used by SNMP for communication. This layer helps in implementing a multi-lingual agent, so that various versions of SNMP can coexist simultaneously in the same network.

The security subsystem features the use of various authentication and privacy protocols with various timeliness checking and engine clock synchronization schemes. SNMPv3 is designed to be secure against:

- Modification of information, where an in-transit message is altered.
- Masquerades, where an unauthorized entity assumes the identity of an authorized entity.
- Message stream modification, where packets are delayed and/or replayed.
- Disclosure, where packet exchanges are sniffed (examined) and information is learned about the contents.

The access control subsystem provides the ability to configure whether access to a managed object in a local MIB is allowed for a remote principal. The access control scheme allows you to define access policies based on MIB views, groups, and multiple security levels.

In addition, the SNMPv3 target and notification MIBs provide a more procedural approach for the generation and filtering of notifications.

SNMPv3 objects are stored in non-volatile memory unless specifically assigned to volatile storage. Objects defined as permanent cannot be deleted or modified.

NOTE

In SNMPv3, many objects can be identified by a human-readable string or by a string of hex octets. In many commands, you can use either a character string, or a colon separated string of hex octets to specify objects. This is indicated by the keyword `hex` *used in the command.*

Message Processing

A particular network manager may require messages that conform to a particular version of SNMP. The choice of the SNMPv1, SNMPv2, or SNMPv3 message processing model can be configured for each network manager as its target address is configured. The selection of the message processing model is configured with the `mp-model` keyword in the following command:

```
configure snmpv3 add target-params {hex} <param name> user {hex} <user name> mp-model
[snmpv1 | snmpv2c | snmpv3] sec-model [snmpv1 | snmpv2c | usm] {sec-level [noauth |
authnopriv | priv]} {volatile}
```

SNMPv3 Security

In SNMPv3 the User-Based Security Model (USM) for SNMP was introduced. USM deals with security related aspects like authentication, encryption of SNMP messages and defining users and their various access security levels. This standard also encompass protection against message delay and message replay.

USM Timeliness Mechanisms

There is one SNMPv3 engine on an Extreme switch, identified by its *snmpEngineID*. The first four octets are fixed to 80:00:07:7C, which represents the Extreme Networks Vendor ID. By default, the additional octets for the snmpEngineID are generated from the device MAC address. Every SNMPv3 engine necessarily maintains two objects: *SNMPEngineBoots*, which is the number of reboots the agent has experienced and *SNMPEngineTime*, which is the engine local time since reboot. It has a local copy of these objects and the *latestReceivedEngineTime* for every authoritative engine it wants to communicate with. Comparing these objects with the values received in messages and then applying certain rules to decide upon the message validity accomplish protection against message delay or message replay.

In a chassis, the `snmpEngineID` will be generated using the MAC address of the MSM with which the switch boots first. For MSM hitless failover, the same `snmpEngineID` will be propagated to both of the MSMs.

The *snmpEngineID* can be configured from the command line, but once the `snmpEngineID` is changed, default users will be reverted back to their original passwords/keys, while non-default users will be reset to the security level of no authorization, no privacy. Use the following command to set the *snmpEngineID*:

```
configure snmpv3 engine-id <hex octet>
```

SNMPEngineBoots can also be configured from the command line. *SNMPEngineBoots* can be set to any desired value but will latch on its maximum, 2147483647. Use the following command to set the *SNMPEngineBoots*:

```
configure snmpv3 engine-boots <(1-2147483647)>
```

Users, Groups, and Security

SNMPv3 controls access and security using the concepts of users, groups, security models, and security levels.

Users. Users are created by specifying a user name. Depending on whether the user will be using authentication and/or privacy, you would also specify an authentication protocol (MD5 or SHA) with password or key, and/or privacy (DES) password or key. To create a user, use the following command:

```
configure snmpv3 add user {hex} <user name> {authentication [md5 | sha] [hex <hex
octet> | <password>]} {privacy [hex <hex octet> | <password>]} {volatile}
```

There are a number of default, permanent users initially available.The default user names are: *admin, initial, initialmd5, initialsha, initialmd5Priv, initialshaPriv.* The default password for *admin* is *password*. For the other default users, the default password is the user name.

To display information about a user, or all users, use the following command:

```
show snmpv3 user {{hex} <user name>}
```

To delete a user, use the following command:

```
configure snmpv3 delete user [all-non-defaults | {hex} <user name>]
```

NOTE

In the SNMPv3 specifications there is the concept of a security name. In the ExtremeWare implementation, the user name and security name are identical. In this manual we use both terms to refer to the same thing.

Groups. Groups are used to manage access for the MIB. You use groups to define the security model, the security level, and the portion of the MIB that members of the group can read or write. To underscore the access function of groups, groups are defined using the following command:

```
configure snmpv3 add access {hex} <group name> {sec-model [snmpv1 | snmpv2 | usm]}
{sec-level [noauth | authnopriv | authpriv]} {read-view {hex} <view name>} {
write-view {hex} <view name>} {notify-view {hex} <view name>} {volatile}
```

The security model and security level are discussed in the section labeled “Security Models and Levels”. The view names associated with a group define a subset of the MIB (subtree) that can be accessed by members of the group. The read view defines the subtree that can be read, write view defines the subtree that can be written to, and notify view defines the subtree that notifications can originate from. MIB views are discussed in the section “MIB Access Control”.

There are a number of default (permanent) groups already defined. These groups are: *admin, initial, initialmd5, initialsha, initialmd5Priv, initialshaPriv, v1v2c_ro, v1v2c_rw.* Use the following command to display information about the access configuration of a group or all groups:

```
show snmpv3 access {{hex} <group name>}
```

Users are associated with groups using the following command:

```
configure snmpv3 add group {hex} <group name> user {hex} <user name> {sec-model
[snmpv1| snmpv2 | usm]} {volatile}
```

To show which users are associated with a group, use the following command:

```
show snmpv3 group {{hex} <group name> {user {hex} <user name>}}
```

To delete a group, use the following command:

```
configure snmpv3 delete access [all-non-defaults | {{hex} <group name>
{sec-model [snmpv1 | snmpv2c | usm] sec-level [noauth | authnopriv |
priv]}}]
```

When you delete a group, you do not remove the association between the group. To delete the association between a user and a group, use the following command:

```
configure snmpv3 delete group {{hex} <group name>} user [all-non-defaults | {{hex}
<user name> {sec-model [snmpv1|snmpv2c|usm]}}]
```

Security Models and Levels. For compatibility, SNMPv3 supports three security models:

- SNMPv1—no security

- SNMPv2c—community strings based security
- SNMPv3—USM security

The default is User-Based Security Model (USM). You can select the security model based on the network manager in your network.

The three security levels supported by USM are:

- noAuthnoPriv—No authentication, no privacy. This is the case with existing SNMPv1/v2c agents.
- AuthnoPriv—Authentication, no privacy. Messages are tested only for authentication.
- AuthPriv—Authentication, privacy. This represents the highest level of security and requires every message exchange to pass the authentication and encryption tests.

When a user is created, an authentication method is selected, and the authentication and privacy passwords or keys are entered.

When MD5 authentication is specified, HMAC-MD5-96 is used to achieve authentication with a 16-octet key, which generates an 128-bit authorization code. This code is inserted in msgAuthenticationParameters field of SNMPv3 PDUs when the security level is specified as either AuthnoPriv or AuthPriv. Specifying SHA authentication uses the HMAC-SHA protocol with a 20-octet key for authentication.

For privacy, a 16-octet key is provided as input to DES-CBS encryption protocol, which generates an encrypted PDU to be transmitted. DES uses bytes 1-7 to make a 56 bit key. This key (encrypted itself) is placed in msgPrivacyParameters of SNMPv3 PDUs when the security level is specified as AuthPriv.

MIB Access Control

SNMPv3 provides a fine-grained mechanism for defining which parts of the MIB can be accessed. This is referred to as the View-Based Access Control Model (VACM).

MIB views represent the basic building blocks of VACM. They are used to define a subset of the information in the MIB. Access to read, to write, and to generate notifications is based on the relationship between a MIB view and an access group. The users of the access group can then read, write, or receive notifications from the part of the MIB defined in the MIB view as configured in the access group.

A view name, a MIB subtree/mask, and an inclusion or exclusion define every MIB view. For example, there is a *System* group defined under the MIB-2 tree. The Object Identifier (OID) for MIB-2 is 1.3.6.1.2, and the *System* group is defined as MIB-2.1.1, or directly as 1.3.6.1.2.1.1.

To define a MIB view which includes only the *System* group, use the following subtree/mask combination:

```
1.3.6.1.2.1.1 / 1.1.1.1.1.1.1.0
```

The mask can also be expressed in hex notation (this is used for the ExtremeWare CLI):

```
1.3.6.1.2.1.1 / fe
```

To define a view that includes the entire MIB-2, use the following subtree/mask:

```
1.3.6.1.2.1.1 / 1.1.1.1.1.0.0.0
```

which, on the command line, is:

```
1.3.6.1.2.1.1 / f8
```

When you create the MIB view, you can choose to include the MIB subtree/mask, or to exclude the MIB subtree/mask. To create a MIB view, use the following command:

```
configure snmpv3 add mib-view {hex} <view name> subtree <object identifier> {/<subtree
mask>} {type [included | excluded]} {volatile}
```

Once the view is created, you can repeatedly use the `configure snmpv3 add mib-view` command to include and/or exclude MIB subtree/mask combinations to precisely define the items you wish to control access to.

In addition to the user created MIB views, there are three default views. They are of storage type permanent and cannot be deleted, but they can be modified. The default views are: *defaultUserView*, *defaultAdminView*, and *defaultNotifyView*. To show MIB views, use the following command:

```
show snmpv3 mib-view {{hex} <view name> {subtree <object identifier>}}
```

To delete a MIB view, use the following command:

```
configure snmpv3 delete mib-view [all-non-defaults | {{hex} <view name> {subtree
<object identifier>}}]
```

MIB views which are being used by security groups cannot be deleted.

Notification

SNMPv3 notification is an enhancement to the concept of SNMP traps. Notifications are messages sent from an agent to the network manager, typically in response to some state change on the agent system. With SNMPv3, you can define precisely which traps you want sent, to which receiver by defining filter profiles to use for the notification receivers.

To configure notifications, you will configure a target address for the process that receives the notification, a target parameters name, and a list of notification tags. The target parameters specify the security and message processing models to use for the notifications to the target. The target parameters name also points to the filter profile used to filter the notifications. Finally, the notification tags are added to a notification table so that any target addresses using that tag will receive notifications.

Target Addresses

A target address is similar to the earlier concept of a trap receiver. To configure a target address, use the following command:

```
configure snmpv3 add target-addr {hex} <addr name> param {hex} <param name> ipaddress
<ip address> {transport-port <port>} {from <source IP address>} {tag-list {hex} <tag>,
{hex} <tag>, ...} {volatile}
```

In configuring the target address you will supply an address name that will be used to identify the target address, a parameters name that will indicate the message processing model and security for the messages sent to the target address, and the IP address and port for the receiver. The parameters name also is used to indicate the filter profile used for notifications. The target parameters are discussed in the section “Target Parameters” on page 58.

The `from` option sets the source IP address in the notification packets.

The `tag-list` option allows you to associate a list of tags with the target address. The tag *defaultNotify* is set by default. Tags are discussed in the section "Notification Tags".

To display target addresses, use the following command:

```
show snmpv3 target-addr {{hex} <addr name>}
```

To delete a single target address or all target addresses, use the following command:

```
configure snmpv3 delete target-addr [{{hex} <addr name>} | all]
```

Target Parameters

Target parameters specify the message processing model, security model, security level, and user name (security name) used for messages sent to the target address. See the sections "Message Processing" on page 53 and "Users, Groups, and Security" on page 54 for more details on these topics. In addition, the target parameter name used for a target address points to a filter profile used to filter notifications. When you specify a filter profile, you associate it with a parameter name, so you need to create different target parameter names if you use different filters for different target addresses.

Use the following command to create a target parameter name, and set the message processing and security settings associated with it:

```
configure snmpv3 add target-params {hex} <param name> user {hex} <user name> mp-model
[snmpv1 | snmpv2c | snmpv3] sec-model [snmpv1 | snmpv2c | usm] {sec-level [noauth |
authnopriv | priv]} {volatile}
```

To display the options associated with a target parameters name, or all target parameters names, use the following command:

```
show snmpv3 target-params {{hex} <param name>}
```

To delete one or all the target parameters, use the following command:

```
configure snmpv3 delete target-params [{{hex} <param name>} | all]
```

Filter Profiles and Filters

A filter profile is a collection of filters that specifies which notifications should be sent to a target address. A filter is defined by a MIB subtree and mask, and by whether that subtree and mask is included or excluded from notification.

When you create a filter profile, you are only associating a filter profile name with a target parameter name. The filters that make up the profile are created and associated with the profile using a different command. To create a filter profile, use the following command:

```
configure snmpv3 add filter-profile {hex} <profile name> param {hex} <param name>
{volatile}
```

Once the profile name is created, you can associate filters with it using the following command:

```
configure snmpv3 add filter {hex} <profile name> subtree <object identifier>
{/<subtree mask>} type [included | excluded] {volatile}
```

The MIB subtree and mask are discussed in the section "MIB Access Control" on page 56, as filters are closely related to MIB views. You can add filters together, including and excluding different subtrees of the MIB until your filter meets your needs.

To display the association between parameter names and filter profiles, use the following command:

```
show snmpv3 filter-profile {{hex} <profile name>} {param {hex} <param name>}
```

To display the filters that belong a filter profile, use the following command:

```
show snmpv3 filter {{hex} <profile name> {{subtree} <object identifier>}
```

To delete a filter or all filters from a filter profile, use the following command:

```
configure snmpv3 delete filter [all | [{hex} <profile name> {subtree <object
identifier>}]]
```

To remove the association of a filter profile or all filter profiles with a parameter name, use the following command:

```
configure snmpv3 delete filter-profile [all |[{hex}<profile name> {param {hex}<param
name>}]]
```

Notification Tags

When you create a target address, you associate a list of notification tags with the target, or by default, the *defaultNotify* tag is associated with the target. When notifications are generated, only targets associated with tags currently in an internal structure, called *snmpNotifyTable*, will be notified. To add an entry to the table, use the following command:

```
configure snmpv3 add notify {hex} <notify name> tag {hex} <tag> {volatile}
```

Any targets associated with tags in the *snmpNotifyTable* will be notified, based on the filter profile associated with the target.

To display the notifications that are set, use the following command:

```
show snmpv3 notify {{hex} <notify name>}
```

To delete an entry from the *snmpNotifyTable*, use the following command:

```
configure snmpv3 delete notify [{{hex} <notify name>} | all-non-defaults]
```

You cannot delete the default entry from the table, so any targets configured with the *defaultNotify* tag will always receive notifications consistent with any filter profile specified.

Configuring Notifications

Since the target parameters name is used to point to a number of objects used for notifications, configure the target parameter name entry first. You can then configure the target address, filter profiles and filters, and any necessary notification tags.

Authenticating Users

ExtremeWare provides two methods to authenticate users who login to the switch:

- RADIUS client
- TACACS+

NOTE

You cannot configure RADIUS and TACACS+ at the same time.

RADIUS Client

Remote Authentication Dial In User Service (RADIUS, RFC 2138) is a mechanism for authenticating and centrally administrating access to network nodes. The ExtremeWare RADIUS client implementation allows authentication for Telnet, Vista, or console access to the switch.

TACACS+

Terminal Access Controller Access Control System Plus (TACACS+) is a mechanism for providing authentication, authorization, and accounting on a centralized server, similar in function to the RADIUS client. The ExtremeWare version of TACACS+ is used to authenticate prospective users who are attempting to administer the switch. TACACS+ is used to communicate between the switch and an authentication database.

Configuring RADIUS Client and TACACS+

For detailed information about configuring a RADIUS client or TACACS+, see Chapter 9.

Using Network Login

Network login is a feature designed to control the admission of user packets into a network by giving addresses only to users that have been properly authenticated. Network login is controlled by an administrator on a per port, per VLAN basis and uses an integration of DHCP, user authentication over the web interface, and, sometimes, a RADIUS server to provide a user database or specific configuration details.

When network login is enabled on a port in a VLAN, that port will not forward any packets until authentication takes place.

For detailed information about using Network login, see Chapter 9.

Using the Simple Network Time Protocol

ExtremeWare supports the client portion of the Simple Network Time Protocol (SNTP) Version 3 based on RFC1769. SNTP can be used by the switch to update and synchronize its internal clock from a Network Time Protocol (NTP) server. When enabled, the switch sends out a periodic query to the indicated NTP server, or the switch listens to broadcast NTP updates. In addition, the switch supports the configured setting for Greenwich Mean time (GMT) offset and the use of Daylight Saving Time. These features have been tested for year 2000 compliance.

Configuring and Using SNTP

To use SNTP, follow these steps:

1. Identify the host(s) that are configured as NTP server(s). Additionally, identify the preferred method for obtaining NTP updates. The options are for the NTP server to send out broadcasts, or for switches using NTP to query the NTP server(s) directly. A combination of both methods is possible. You must identify the method that should be used for the switch being configured.
2. Configure the Greenwich Mean Time (GMT) offset and Daylight Saving Time preference. The command syntax to configure GMT offset and usage of Daylight Saving Time is as follows:

   ```
   configure timezone {name <std_timezone_ID>} <GMT_offset> {autodst {name
   <dst_timezone_ID>} {<dst_offset>} {begins [every <floatingday> | on <absoluteday>]
   {at <time_of_day>} {ends [every <floatingday> | on <absoluteday>] {at
   <time_of_day>}}} | noautodst}
   ```

   By default, Daylight Saving Time is assumed to begin on the first Sunday in April at 2:00 AM, and end the last Sunday in October at 2:00 AM, and be offset from standard time by one hour. If this is the case in your timezone, you can set up automatic daylight savings adjustment with the command:

   ```
   configure timezone <GMT_offset> autodst
   ```

   If your timezone uses starting and ending dates and times that differ from the default, you can specify the starting and ending date and time in terms of a floating day, as follows:

   ```
   configure timezone name MET 60 autodst name MDT begins every last sunday march at
   1 ends every last sunday october at 1
   ```

   You can also specify a specific date and time, as shown in the following command.

   ```
   configure timezone name NZST 720 autodst name NZDT 60 begins every first sunday
   october at 2 ends on 3/16/2002 at 2
   ```

   The optional timezone IDs are used to identify the timezone in display commands such as `show switch`.

   Table 12 describes the command options in detail:

Table 12: Time Zone Configuration Command Options

GMT_offset	Specifies a Greenwich Mean Time (GMT) offset, in + or - minutes.
std-timezone-ID	Specifies an optional name for this timezone specification. May be up to six characters in length. The default is an empty string.
autodst	Enables automatic Daylight Savings Time.
dst-timezone-ID	Specifies an optional name for this DST specification. May be up to six characters in length. The default is an empty string.
dst_offset	Specifies an offset from standard time, in minutes. Value is in the range of 1 to 60. Default is 60 minutes.
floating_day	Specifies the day, week, and month of the year to begin or end DST each year. Format is: <week><day><month> where: • <week> is specified as [first | second | third | fourth | last] or 1-5 • <day> is specified as [sunday | monday | tuesday | wednesday | thursday | friday | saturday] or 1-7 (where 1 is Sunday) • <month> is specified as [january | february | march | april | may | june | july | august | september | october | november | december] or 1-12 Default for beginning is first sunday april; default for ending is last sunday october.

Table 12: Time Zone Configuration Command Options (Continued)

absolute_day	Specifies a specific day of a specific year on which to begin or end DST. Format is: <month>/<day>/<year> where: • <month> is specified as 1-12 • <day> is specified as 1-31 • <year> is specified as 1970 - 2035 The year must be the same for the begin and end dates.
time_of_day	Specifies the time of day to begin or end Daylight Savings Time. May be specified as an hour (0-23) or as hour:minutes. Default is 2:00.
noautodst	Disables automatic Daylight Savings Time.

Automatic Daylight Savings Time (DST) changes can be enabled or disabled. The default setting is enabled. To disable automatic DST, use the command:

```
configure timezone {name <std_timezone_ID>} <GMT_offset> noautodst
```

3 Enable the SNTP client using the following command:

```
enable sntp-client
```

Once enabled, the switch sends out a periodic query to the NTP servers defined later (if configured) or listens to broadcast NTP updates from the network. The network time information is automatically saved into the on-board real-time clock.

4 If you would like this switch to use a directed query to the NTP server, configure the switch to use the NTP server(s). If the switch listens to NTP broadcasts, skip this step. To configure the switch to use a directed query, use the following command:

```
configure sntp-client [primary | secondary] server <host name/ip>]
```

NTP queries are first sent to the primary server. If the primary server does not respond within 1 second, or if it is not synchronized, the switch queries the secondary server (if one is configured). If the switch cannot obtain the time, it restarts the query process. Otherwise, the switch waits for the `sntp-client update interval` before querying again.

5 Optionally, the interval for which the SNTP client updates the real-time clock of the switch can be changed using the following command:

```
configure sntp-client update-interval <seconds>
```

The default `sntp-client update-interval` value is 64 seconds.

6 You can verify the configuration using the following commands:

– `show sntp-client`

This command provides configuration and statistics associated with SNTP and its connectivity to the NTP server.

– `show switch`

This command indicates the GMT offset, the Daylight Savings Time configuration and status, and the current local time.

NTP updates are distributed using GMT time. To properly display the local time in logs and other timestamp information, the switch should be configured with the appropriate offset to GMT based on geographical location. Table 13 describes GMT offsets.

Table 13: Greenwich Mean Time Offsets

GMT Offset in Hours	GMT Offset in Minutes	Common Time Zone References	Cities
+0:00	+0	GMT - Greenwich Mean	London, England; Dublin, Ireland; Edinburgh, Scotland; Lisbon, Portugal; Reykjavik, Iceland; Casablanca, Morocco
		UT or UTC - Universal (Coordinated)	
		WET - Western European	
-1:00	-60	WAT - West Africa	Azores, Cape Verde Islands
-2:00	-120	AT - Azores	
-3:00	-180		Brasilia, Brazil; Buenos Aires, Argentina; Georgetown, Guyana;
-4:00	-240	AST - Atlantic Standard	Caracas; La Paz
-5:00	-300	EST - Eastern Standard	Bogota, Columbia; Lima, Peru; New York, NY, Trevor City, MI USA
-6:00	-360	CST - Central Standard	Mexico City, Mexico
-7:00	-420	MST - Mountain Standard	Saskatchewan, Canada
-8:00	-480	PST - Pacific Standard	Los Angeles, CA, Cupertino, CA, Seattle, WA USA
-9:00	-540	YST - Yukon Standard	
-10:00	-600	AHST - Alaska-Hawaii Standard	
		CAT - Central Alaska	
		HST - Hawaii Standard	
-11:00	-660	NT - Nome	
-12:00	-720	IDLW - International Date Line West	
+1:00	+60	CET - Central European	Paris, France; Berlin, Germany; Amsterdam, The Netherlands; Brussels, Belgium; Vienna, Austria; Madrid, Spain; Rome, Italy; Bern, Switzerland; Stockholm, Sweden; Oslo, Norway
		FWT - French Winter	
		MET - Middle European	
		MEWT - Middle European Winter	
		SWT - Swedish Winter	
+2:00	+120	EET - Eastern European, Russia Zone 1	Athens, Greece; Helsinki, Finland; Istanbul, Turkey; Jerusalem, Israel; Harare, Zimbabwe
+3:00	+180	BT - Baghdad, Russia Zone 2	Kuwait; Nairobi, Kenya; Riyadh, Saudi Arabia; Moscow, Russia; Tehran, Iran
+4:00	+240	ZP4 - Russia Zone 3	Abu Dhabi, UAE; Muscat; Tblisi; Volgograd; Kabul
+5:00	+300	ZP5 - Russia Zone 4	
+5:30	+330	IST – India Standard Time	New Delhi, Pune, Allahabad, India
+6:00	+360	ZP6 - Russia Zone 5	
+7:00	+420	WAST - West Australian Standard	
+8:00	+480	CCT - China Coast, Russia Zone 7	
+9:00	+540	JST - Japan Standard, Russia Zone 8	
+10:00	+600	EAST - East Australian Standard	
		GST - Guam Standard	
		Russia Zone 9	

Table 13: Greenwich Mean Time Offsets (Continued)

GMT Offset in Hours	GMT Offset in Minutes	Common Time Zone References	Cities
+11:00	+660		
+12:00	+720	IDLE - International Date Line East NZST - New Zealand Standard NZT - New Zealand	Wellington, New Zealand; Fiji, Marshall Islands

SNTP Example

In this example, the switch queries a specific NTP server and a backup NTP server. The switch is located in Cupertino, CA, and an update occurs every 20 minutes. The commands to configure the switch are as follows:

```
configure timezone -480 autodst
configure sntp-client update interval 1200
enable sntp-client
configure sntp-client primary server 10.0.1.1
configure sntp-client secondary server 10.0.1.2
```

3 Accessing the Switch

This chapter covers the following topics:

Understanding the Command Syntax

This section describes the steps to take when entering a command. Refer to the sections that follow for detailed information on using the command line interface.

ExtremeWare command syntax is described in detail in the *ExtremeWare 7.2e Command Reference Guide.* Some commands are also described in this user guide, in order to describe how to use the features of the ExtremeWare software. However, only a subset of commands are described here, and in some cases only a subset of the options that a command supports. The *ExtremeWare 7.2e Command Reference Guide* should be considered the definitive source for information on ExtremeWare commands.

When entering a command at the prompt, ensure that you have the appropriate privilege level. Most configuration commands require you to have the administrator privilege level. To use the command line interface (CLI), follow these steps:

1 Enter the command name.

If the command does not include a parameter or values, skip to step 3. If the command requires more information, continue to step 2.

2 If the command includes a parameter, enter the parameter name and values.

3 The value part of the command specifies how you want the parameter to be set. Values include numerics, strings, or addresses, depending on the parameter.

4 After entering the complete command, press [Return].

NOTE

If an asterisk (*) appears in front of the command-line prompt, it indicates that you have outstanding configuration changes that have not been saved. For more information on saving configuration changes, see Appendix B.

Syntax Helper

The CLI has a built-in syntax helper. If you are unsure of the complete syntax for a particular command, enter as much of the command as possible and press [Tab]. The syntax helper provides a list of options for the remainder of the command, and places the cursor at the end of the command you have entered so far, ready for the next option.

If the command is one where the next option is a named component, such as a VLAN, access profile, or route map, the syntax helper will also list any currently configured names that might be used as the next option. In situations where this list might be very long, the syntax helper will list only one line of names, followed by an ellipses to indicate that there are more names than can be displayed.

The syntax helper also provides assistance if you have entered an incorrect command.

Abbreviated Syntax

Abbreviated syntax is the shortest unambiguous allowable abbreviation of a command or parameter. Typically, this is the first three letters of the command. If you do not enter enough letters to allow the switch to determine which command you mean, the syntax helper will provide a list of the options based on the portion of the command you have entered.

NOTE

When using abbreviated syntax, you must enter enough characters to make the command unambiguous and distinguishable to the switch.

Command Shortcuts

All named components of the switch configuration must have a unique name. Components are typically named using the `create` command. When you enter a command to configure a named component, you do not need to use the keyword of the component. For example, to create a VLAN, you must enter a unique VLAN name:

```
create vlan engineering
```

Once you have created the VLAN with a unique name, you can then eliminate the keyword `vlan` from all other commands that require the name to be entered. For example, instead of entering the switch command:

```
configure vlan engineering delete port 1-3,6
```

you could enter the following shortcut:

```
configure engineering delete port 1-3,6
```

Switch Numerical Ranges

Commands that require you to enter one or more port numbers use the parameter `<portlist>` in the syntax. A portlist can be a range of numbers, for example:

```
port 1-3
```

You can add additional port numbers to the list, separated by a comma:

```
port 1-3,6,8
```

Names

All named components of the switch configuration must have a unique name. Names must begin with an alphabetical character and are delimited by whitespace, unless enclosed in quotation marks. Names are not case-sensitive. Names cannot be tokens used on the switch.

Symbols

You may see a variety of symbols shown as part of the command syntax. These symbols explain how to enter the command, and you do not type them as part of the command itself. Table 14 summarizes command syntax symbols.

Table 14: Command Syntax Symbols

Symbol	Description
angle brackets < >	Enclose a variable or value. You must specify the variable or value. For example, in the syntax `configure vlan <vlan name> ipaddress <ipaddress>` you must supply a VLAN name for `<vlan name>` and an address for `<ip_address>` when entering the command. Do not type the angle brackets.
square brackets [ ]	Enclose a required value or list of required arguments. One or more values or arguments can be specified. For example, in the syntax `use image [primary | secondary]` you must specify either the primary or secondary image when entering the command. Do not type the square brackets.
vertical bar |	Separates mutually exclusive items in a list, one of which must be entered. For example, in the syntax `configure snmp community [read-only | read-write] <string>` you must specify either the read or write community string in the command. Do not type the vertical bar.
braces { }	Enclose an optional value or a list of optional arguments. One or more values or arguments can be specified. For example, in the syntax `reboot {<date> <time> | cancel}` you can specify either a particular date and time combination, or the keyword `cancel` to cancel a previously scheduled reboot. If you do not specify an argument, the command will prompt, asking if you want to reboot the switch now. Do not type the braces.

Limits

The command line can process up to 200 characters, including spaces. If you enter more than 200 characters, the switch generates a stack overflow error and processes the first 200 characters.

Line-Editing Keys

Table 15 describes the line-editing keys available using the CLI.

Table 15: Line-Editing Keys

Key(s)	Description
Backspace	Deletes character to left of cursor and shifts remainder of line to left.
Delete or [Ctrl] + D	Deletes character under cursor and shifts remainder of line to left.
[Ctrl] + K	Deletes characters from under cursor to end of line.
Insert	Toggles on and off. When toggled on, inserts text and shifts previous text to right.
Left Arrow	Moves cursor to left.
Right Arrow	Moves cursor to right.
Home or [Ctrl] + A	Moves cursor to first character in line.
End or [Ctrl] + E	Moves cursor to last character in line.
[Ctrl] + L	Clears screen and movers cursor to beginning of line.
[Ctrl] + P or Up Arrow	Displays previous command in command history buffer and places cursor at end of command.
[Ctrl] + N or Down Arrow	Displays next command in command history buffer and places cursor at end of command.
[Ctrl] + U	Clears all characters typed from cursor to beginning of line.
[Ctrl] + W	Deletes previous word.

Command History

ExtremeWare “remembers” the last 49 commands you entered. You can display a list of these commands by using the following command:

```
history
```

Common Commands

Table 16 describes some of the common commands used to manage the switch. Commands specific to a particular feature may also be described in other chapters of this guide. For a detailed description of the commands and their options, see the *ExtremeWare 7.2e Command Reference Guide*.

Table 16: Common Commands

Command	Description
clear session <number>	Terminates a Telnet session from the switch.

Table 16: Common Commands (Continued)

Command	Description
configure account <user account> {encrypted} {<password>}	Configures a user account password. The switch will interactively prompt for a new password, and for reentry of the password to verify it. Passwords must have a minimum of 1 character and can have a maximum of 30 characters. Passwords are case-sensitive; user names are not case sensitive.
configure banner	Configures the banner string. You can enter up to 24 rows of 79-column text that is displayed before the login prompt of each session. Press [Return] at the beginning of a line to terminate the command and apply the banner. To clear the banner, press [Return] at the beginning of the first line.
configure banner netlogin	Configures the network login banner string. You can enter up to 1024 characters to be displayed before the login prompt of each session.
configure ports [<portlist> | all | mgmt] auto off {speed [10 | 100 | 1000]} duplex [half | full]	Manually configures the port speed and duplex setting of one or more ports on a switch.
configure ssh2 key {pregenerated}	Generates the SSH2 host key.
configure sys-recovery-level [none | [all | critical] [reboot | shutdown | system-dump [maintenance-mode | reboot | shutdown]]]	Configures a recovery option for instances where an exception occurs in ExtremeWare.
configure time <date> <time>	Configures the system date and time. The format is as follows: `mm/dd/yyyy hh:mm:ss` The time uses a 24-hour clock format. You cannot set the year past 2036.
configure timezone {name <std_timezone_ID>} <GMT_offset> {autodst {name <dst_timezone_ID>} {<dst_offset>} {begins [every <floatingday> | on <absoluteday>] {at <time_of_day>} {ends [every <floatingday> | on <absoluteday>] {at <time_of_day>}}} | noautodst}	Configures the time zone information to the configured offset from GMT time. The format of `gmt_offset` is +/- minutes from GMT time. The `autodst` and `noautodst` options enable and disable automatic Daylight Saving Time change based on the North American standard. Additional options are described in the *ExtremeWare 7.2e Command Reference Guide*.
configure vlan <vlan name> ipaddress <ipaddress> {<netmask> | <mask length>}	Configures an IP address and subnet mask for a VLAN.
create account [admin | user] <username> {encrypted} {<password>}	Creates a user account. This command is available to admin-level users and to users with RADIUS command authorization. The username is between 1 and 30 characters, the password is between 0 and 30 characters.
create vlan <vlan name>	Creates a VLAN.
delete account <username>	Deletes a user account.
delete vlan <vlan name>	Deletes a VLAN.
disable bootp vlan [<vlan name> | all]	Disables BOOTP for one or more VLANs.
disable cli-config-logging	Disables logging of CLI commands to the Syslog.
disable clipaging	Disables pausing of the screen display when a show command output reaches the end of the page.
disable idletimeouts	Disables the timer that disconnects all sessions. Once disabled, console sessions remain open until the switch is rebooted or you logoff. Telnet sessions remain open until you close the Telnet client.
disable ports [<portlist> | all]	Disables a port on the switch.

Table 16: Common Commands (Continued)

Command	Description
disable ssh2	Disables SSH2 Telnet access to the switch.
disable telnet	Disables Telnet access to the switch.
disable web	Disables web access to the switch.
enable bootp vlan [<vlan name> | all]	Enables BOOTP for one or more VLANs.
enable cli-config-logging	Enables the logging of CLI configuration commands to the Syslog for auditing purposes. The default setting is enabled.
enable clipaging	Enables pausing of the screen display when `show command` output reaches the end of the page. The default setting is enabled.
enable idletimeouts	Enables a timer that disconnects all sessions (both Telnet and console) after 20 minutes of inactivity. The default setting is disabled.
enable license [basic_L3 | advanced_L3 | full_L3] <license_key>	Enables a particular software feature license. Specify `<license_key>` as an integer. The command `unconfigure switch {all}` does not clear licensing information. This license cannot be disabled once it is enabled on the switch.
enable ssh2 {access-profile [<access profile> | none]} {port <tcp_port_number>}	Enables SSH2 sessions. By default, SSH2 is enabled with no access profile, and uses TCP port number 22. To cancel a previously configured access-profile, use the `none` option.
enable telnet {access-profile [<access_profile> | none]} {port <tcp_port_number>}	Enables Telnet access to the switch. By default, Telnet is enabled with no access profile, and uses TCP port number 23. To cancel a previously configured access-profile, use the `none` option.
enable web {access-profile [<access_profile> | none]} {port <tcp_port_number>}	Enables ExtremeWare Vista™ web access to the switch. By default, web access is enabled with no access profile, using TCP port number 80. Use the `none` option to cancel a previously configured access-profile.
history	Displays the previous 49 commands entered on the switch.
show banner	Displays the user-configured banner.
unconfigure switch {all}	Resets all switch parameters (with the exception of defined user accounts, and date and time information) to the factory defaults. If you specify the keyword `all`, the switch erases the currently selected configuration image in flash memory and reboots. As a result, all parameters are reset to default settings.

Configuring Management Access

ExtremeWare supports the following two levels of management:

- User
- Administrator

In addition to the management levels, you can optionally use an external RADIUS server to provide CLI command authorization checking for each command. For more information on RADIUS, see “RADIUS Client” in Chapter 2.

User Account

A user-level account has viewing access to all manageable parameters, with the exception of:

- User account database.
- SNMP community strings.

A user-level account can use the `ping` command to test device reachability, and change the password assigned to the account name. If you have logged on with user capabilities, the command-line prompt ends with a (>) sign. For example:

```
Summit2>
```

Administrator Account

An administrator-level account can view and change all switch parameters. It can also add and delete users, and change the password associated with any account name. The administrator can disconnect a management session that has been established by way of a Telnet connection. If this happens, the user logged on by way of the Telnet connection is notified that the session has been terminated.

If you have logged on with administrator capabilities, the command-line prompt ends with a (#) sign. For example:

```
Summit18#
```

Prompt Text

The prompt text is taken from the SNMP `sysname` setting. The number that follows the colon indicates the sequential line/command number.

If an asterisk (*) appears in front of the command-line prompt, it indicates that you have outstanding configuration changes that have not been saved. For example:

```
*Summit9#
```

Default Accounts

By default, the switch is configured with two accounts, as shown in Table 17.

Table 17: Default Accounts

Account Name	Access Level
admin	This user can access and change all manageable parameters. The admin account cannot be deleted.
user	This user can view (but not change) all manageable parameters, with the following exceptions: • This user cannot view the user account database. • This user cannot view the SNMP community strings.

Changing the Default Password

Default accounts do not have passwords assigned to them. Passwords can have a minimum of zero characters and can have a maximum of 30 characters.

NOTE

Passwords are case-sensitive; user names are not case-sensitive.

To add a password to the default admin account, follow these steps:

1. Log in to the switch using the name *admin*.
2. At the password prompt, press [Return].
3. Add a default admin password by entering the following command:

   `configure account admin`
4. Enter the new password at the prompt.
5. Re-enter the new password at the prompt.

To add a password to the default user account, follow these steps:

1. Log in to the switch using the name *admin*.
2. At the password prompt, press [Return], or enter the password that you have configured for the *admin* account.
3. Add a default user password by entering the following command:

   `configure account user`
4. Enter the new password at the prompt.
5. Re-enter the new password at the prompt.

NOTE

If you forget your password while logged out of the command line interface, contact your local technical support representative, who will advise on your next course of action.

Creating a Management Account

The switch can have a total of 16 management accounts. You can use the default names (*admin* and *user*), or you can create new names and passwords for the accounts. Passwords can have a minimum of 0 characters and can have a maximum of 30 characters.

To create a new account, follow these steps:

1. Log in to the switch as *admin*.
2. At the password prompt, press [Return], or enter the password that you have configured for the *admin* account.
3. Add a new user by using the following command:

   `create account [admin | pppuser | user] <username>`
4. Enter the password at the prompt.
5. Re-enter the password at the prompt.

Viewing Accounts

To view the accounts that have been created, you must have administrator privileges. Use the following command to see the accounts:

```
show accounts
```

Deleting an Account

To delete a account, you must have administrator privileges. To delete an account, use the following command:

```
delete account <username>
```

NOTE

Do not delete the default administrator account. If you do, it is automatically restored, with no password, the next time you download a configuration. To ensure security, change the password on the default account, but do not delete it. The changed password will remain intact through configuration uploads and downloads.

If you must *delete the default account, first create another administrator-level account. Remember to manually delete the default account again every time you download a configuration.*

Domain Name Service Client Services

The Domain Name Service (DNS) client in ExtremeWare augments the following commands to allow them to accept either IP addresses or host names:

- `telnet`
- `download [bootrom | configuration | image]`
- `upload configuration`
- `ping`
- `traceroute`

In addition, the `nslookup` utility can be used to return the IP address of a hostname.

You can specify up to eight DNS servers for use by the DNS client using the following command:

```
configure dns-client add <ipaddress>
```

You can specify a default domain for use when a host name is used without a domain. Use the following command:

```
configure dns-client default-domain <domain_name>
```

For example, if you specify the domain “xyz-inc.com” as the default domain, then a command such as `ping accounting1` will be taken as if it had been entered `ping accounting1.xyz-inc.com`.

Checking Basic Connectivity

The switch offers the following commands for checking basic connectivity:

- ping
- traceroute

Ping

The ping command enables you to send Internet Control Message Protocol (ICMP) echo messages to a remote IP device. The ping command is available for both the user and administrator privilege level.

The ping command syntax is:

```
ping {udp} {continuous} {size <start_size> {-<end_size}} [<ip_address> | <hostname>]
{from <src_ipaddress> | with record-route | from <src_ipaddress> with record-route}
```

Options for the ping command are described in Table 18.

Table 18: Ping Command Parameters

Parameter	Description
udp	Specifies that UDP messages should be sent instead of ICMP echo messages. When specified, `from` and `with record-route` options are not supported.
`continuous`	Specifies ICMP echo messages to be sent continuously. This option can be interrupted by pressing any key.
`size`	Specifies the size of the ICMP request. If both the `start_size` and `end_size` are specified, transmits ICMP requests using 1 byte increments, per packet. If no `end_size` is specified, packets of `start_size` are sent.
<ipaddress>	Specifies the IP address of the host.
<hostname>	Specifies the name of the host. To use the `hostname`, you must first configure DNS.
from	Uses the specified source address in the ICMP packet. If not specified, the address of the transmitting interface is used.
with record-route	Decodes the list of recorded routes and displays them when the ICMP echo reply is received.

If a ping request fails, the switch continues to send ping messages until interrupted. Press any key to interrupt a ping request. The statistics are tabulated after the ping is interrupted.

Traceroute

The traceroute command enables you to trace the routed path between the switch and a destination endstation. The traceroute command syntax is:

```
traceroute <host name/ip> {from <source IP address>} {ttl <number>} {port <port
number>}
```

where:

- `ip_address` is the IP address of the destination endstation.
- `hostname` is the hostname of the destination endstation. To use the hostname, you must first configure DNS.

- `from` uses the specified source address in the ICMP packet. If not specified, the address of the transmitting interface is used.
- `ttl` configures the switch to trace the hops until the time-to-live has been exceeded for the switch.
- `port` uses the specified UDP port number.

4 Configuring Ports

This chapter covers the following topics:

Enabling and Disabling Switch Ports

To enable or disable one or more ports, use the following command:

```
enable ports [<portlist> | all]
disable ports [<portlist> | all]
```

For example, to disable ports 3, 5, and 12 through 15, use the following command:

```
disable ports 3,5,12-15
```

Even though a port is disabled, the link remains enabled for diagnostic purposes.

Configuring Switch Port Speed and Duplex Setting

When configuring the speed and duplex setting for a port, autonegotiation plays a significant role. By default, the switch is configured to use autonegotiation to determine the port speed and duplex setting for each port.

The Summit 400-48t has four 1000 Mbps fiber and 48 copper 10/100/1000 Mbps ports. All of the fiber ports can be configured for automatic failover using the first four copper ports. These ports are called *combination* ports because either the fiber port or the copper port is active, but they are never active concurrently. For a description of cabling for combination ports, see “Uplink Redundancy” on page 27. For information on configuring combination ports, see “Configuring Automatic Failover for

Combination Ports” on page 85. If you plan to use the automatic failover feature, ensure that port settings are set correctly for autonegotiation.

Fiber ports run at 1000 Mbps, regardless of whether you attempt to manually slow them using the CLI. If you plan on running the copper ports at 1000 Mbps, it is recommended that you keep autonegotiation on. For ports running at slower speeds, you can manually configure the speed of 10/100/1000 Mbps ports and disable autonegotiation.

10BASE-T and 100BASE-TX ports can connect to either 10BASE-T or 100BASE-T networks. By default, the ports autonegotiate port speed. You can also configure each port for a particular speed (either 10 Mbps or 100 Mbps).

Gigabit Ethernet fiber ports only run at full-duplex but you must specify the `duplex` setting. The 10/100/1000 copper ports, however, can be either be configured for half-duplex or full-duplex operation. These 10/100/1000 Mbps ports are supported only through autonegotiation.

To configure port speed and duplex setting, use the following command:

```
configure ports [<portlist> | all | mgmt] auto off {speed [10 | 100 | 1000]} duplex
[half | full]
```

To configure the system to autonegotiate, use the following command:

```
configure ports [<portlist> | mgmt | all] auto on
```

Summit 400 ports do not advertise or support flow control frames.

Turning Off Autonegotiation for a Gigabit Ethernet Port

In certain interoperability situations, you may need to turn autonegotiation off on a Gigabit Ethernet port. Even though a Gigabit Ethernet port runs only at full duplex, you must specify the duplex setting.

NOTE

1000BASE-T ports support only autonegotiation.

The following example turns autonegotiation off for port 4 (a combination port):

```
configure ports 4 auto off duplex full speed 1000
```

Configuring Link Detection

ExtremeWare contains an interrupt service routine (ISR) that sends interrupts when links transition. If a link continuously transitions, causing the ISR to send continuous interrupts, the middle layer filter filters out the continuous interrupt messages. You can configure the interaction between these functions using the following command:

```
configure ports <portlist> link-detection-level <link-detection-level>
```

Configuring Interpacket Gap for Gigabit Ethernet Ports

On the Summit 400-48t, you can configure the Interpacket Gap for 1 or 10 Gigabit Ethernet ports. The Interpacket Gap, sometimes referred to as the Interframe Gap, is the transmit packet byte-time delay between successive data packets mandated by the IEEE for Ethernet networks. Byte-time is the amount

of time it takes to transmit one byte on the link at the specified or negotiated link speed. The configured Interpacket Gap value has no effect on received packets. The default value is 12. The minimum and maximum allowed values range between 12 and 1023.

The standard effective Interpacket Gap for Gigabit Ethernet interfaces ranges between 12 and 1023. Some vendors' 10 Gigabit Ethernet interfaces drop packets when packets are transmitted using a value of 12. Thus, by increasing the Interpacket Gap, packet transmission is slowed and packet loss can be minimized or prevented. The Interpacket Gap value need not be modified when interconnecting Extreme Networks switches over 10 Gigabit Ethernet links. Use the following command to modify the Interpacket Gap:

```
configure port <port> interpacket-gap <byte_time>
```

Jumbo Frames

Jumbo frames are Ethernet frames that are larger than 1522 bytes, including four bytes used for the cyclic redundancy check (CRC). Extreme products support switching and routing of jumbo frames at wire-speed on all ports.

Jumbo frames are used between endstations that support larger frame sizes for more efficient transfers of bulk data. Both endstations involved in the transfer must be capable of supporting jumbo frames. The switch only performs IP fragmentation, or participates in maximum transmission unit (MTU) negotiation on behalf of devices that support jumbo frames.

Enabling Jumbo Frames

To enable jumbo frame support, enable jumbo frames on the desired ports. To set the maximum jumbo frame size, use the following command:

```
configure jumbo-frame size <number>
```

The jumbo frame size range is 1523 to 9216. This value describes the maximum size of the frame in transit (on the wire), and includes 4 bytes of CRC plus another 4 bytes if 802.1Q tagging is being used.

Set the MTU size for the VLAN, using the following command:

```
configure ip-mtu <number> vlan <vlan name>
```

The IP MTU default is 1500. The range is 1500-9194.

Next, enable support on the physical ports that will carry jumbo frames using the following command:

```
enable jumbo-frame ports [<portlist> | all]
```

NOTE

Some network interface cards (NICs) have a configured maximum MTU size that does not include the additional 4 bytes of CRC. Ensure that the NIC maximum MTU size is at or below the maximum MTU size configured on the switch. Frames that are larger than the MTU size configured on the switch are dropped at the ingress port.

Jumbo Frames Example

The following example create two VLANs *sw1* and *sw2*. It adds port 12 to *sw1* and port 13 to *sw2*. It configures port 12 and 13 for jumbo frames up to 9216 bytes (including CRC). It also configures VLANs *sw1* and *sw2* to accept IP packets up to 9194 bytes.

```
* Summit400-48t:48 # create vlan sw1
* Summit400-48t:49 # create vlan sw2
* Summit400-48t:50 # configure vlan sw1 add port 12
* Summit400-48t:51 # configure vlan sw2 add port 13
* Summit400-48t:52 # configure jumbo-frame size 9216
* Summit400-48t:53 # enable jumbo-frame ports 12,13
* Summit400-48t:54 # configure ip-mtu 9194 vlan vlansw1
* Summit400-48t:55 # configure ip-mtu 9194 vlan vlansw2
```

Path MTU Discovery

Using path MTU discovery, a source host assumes that the path MTU is the MTU of the first hop (which is known). The host sends all datagrams on that path with the “don’t fragment” (DF) bit set, which restricts fragmentation. If any of the datagrams must be fragmented by an Extreme switch along the path, the Extreme switch discards the datagrams and returns an ICMP Destination Unreachable message to the sending host, with a code meaning “fragmentation needed and DF set”. When the source host receives the message (sometimes called a “Datagram Too Big” message), the source host reduces its assumed path MTU and retransmits the datagrams.

The path MTU discovery process ends when one of the following is true:

- The source host sets the path MTU low enough that its datagrams can be delivered without fragmentation.
- The source host does not set the DF bit in the datagram headers.

If it is willing to have datagrams fragmented, a source host can choose not to set the DF bit in datagram headers. Normally, the host continues to set DF in all datagrams, so that if the route changes and the new path MTU is lower, the host can perform path MTU discovery again.

IP Fragmentation with Jumbo Frames

ExtremeWare supports the fragmenting of IP packets. If an IP packet originates in a local network that allows large packets and those packets traverse a network that limits packets to a smaller size, the packets are fragmented instead of discarded.

This feature is designed to be used in conjunction with jumbo frames. Frames that are fragmented are not processed at wire-speed within the switch fabric.

NOTE

Jumbo frame-to-jumbo frame fragmentation is not supported. Only jumbo frame-to-normal frame fragmentation is supported.

To configure VLANs for IP fragmentation, follow these steps:

1. Enable jumbo frames on the incoming port.
2. Add the port to a VLAN.

3 Assign an IP address to the VLAN.

4 Enable ipforwarding on the VLAN.

5 Set the MTU size for the VLAN, using the following command:

```
configure ip-mtu <number> vlan <vlan name>
```

The ip-mtu value can be 1500 - 9194, with 1500 the default.

NOTE

To set the MTU size greater than 1500, all ports in the VLAN must have jumbo frames enabled.

IP Fragmentation within a VLAN

ExtremeWare supports IP fragmentation within a VLAN. This feature does not require you to configure the MTU size. To use IP fragmentation within a VLAN, follow these steps:

1 Enable jumbo frames on the incoming port.

2 Add the port to a VLAN.

3 Assign an IP address to the VLAN.

4 Enable ipforwarding on the VLAN.

If you leave the MTU size configured to the default value, when you enable jumbo frame support on a port on the VLAN you receive a warning that the *ip-mtu* size for the VLAN is not set at maximum jumbo frame size. You can ignore this warning if you want IP fragmentation within the VLAN, only.

However, if you do not use jumbo frames, IP fragmentation can only be used for traffic that stays within the same VLAN. To use IP fragmentation for traffic that is set to other VLANs, you must configure all ports in the VLAN for jumbo frame support.

Load Sharing on the Switch

Load sharing allows you to increase bandwidth and availability by using a group of ports to carry traffic in parallel between switches. Load sharing allows the switch to use multiple ports as a single logical port. For example, VLANs see the load-sharing group as a single logical port. Most load-sharing algorithms guarantee packet sequencing between clients.

If a port in a load-sharing group fails, traffic is redistributed to the remaining ports in the load-sharing group. If the failed port becomes active again, traffic is redistributed to include that port.

NOTE

Load sharing must be enabled on both ends of the link or a network loop may result. The load-sharing algorithms do not need to be the same on both ends.

Static Load Sharing

Static load sharing is a grouping of ports specifically configured to load share. The switch ports at each end must be configured as part of a load-sharing group. Additionally, you can choose the load-sharing algorithm used by the group. This feature is supported between Extreme Networks switches only, but

may be compatible with third-party trunking or link-aggregation algorithms. Check with an Extreme Networks technical representative for more information.

Load-Sharing Algorithm

The Summit 400-48 uses an address-based load-sharing algorithm as the distribution technique to determine the output port selection. Algorithm selection is not intended for use in predictive traffic engineering. The algorithm uses addressing information to determine which physical port in the load-sharing group to use for forwarding traffic out of the switch. Addressing information is based on the packet protocol, as follows:

— **IP packets**—Uses the source and destination MAC and IP addresses.

— **All other packets**—Uses the source and destination MAC address.

Configured IP Address-Based Load Sharing

When you configure load sharing, the switch examines a specific place in the packet to determine which egress port to use for forwarding traffic.

You can control the field examined by the switch for IP address-based load sharing, using the following command:

```
configure sharing address-based [ip-dest| ip-source| ip-source-dest |mac-dest |
mac-source | mac-source-dest]
```

where:

ip-dest	Indicates that the switch should examine the IP destination address.
ip-source	Indicates that the switch should examine the IP source address.
ip-source-dest	Indicates that the switch should examine the IP source and destination addresses.
mac-dest	Indicates that the switch should examine the MAC destination address.
mac-source	Indicates that the switch should examine the MAC source address.
mac-source-dest	Indicates that the switch should examine the MAC source and destination addresses.

- ip-dest—Indicates that the switch should examine the IP destination address.
- ip-source—Indicates that the switch should examine the IP source address.
- ip-source-dest—Indicates that the switch should examine the IP source and destination addresses.
- mac-dest—Indicates that the switch should examine the MAC destination address.
- mac-source—Indicates that the switch should examine the MAC source address.
- mac-source-dest—Indicaes that the switch should examine the MAC source and destination addresses.

—This feature is available for the address-based load-sharing algorithm, only.

To verify your configuration, use the following command:

```
show sharing address-based
```

Configuring Switch Load Sharing

To set up a switch to load share among ports, you must create a load-sharing group of ports. The first port in the load-sharing group is configured to be the “master” logical port. This is the reference port used in configuration commands. It can be thought of as the logical port representing the entire port group.

All the ports in a load-sharing group must have the same exact configuration, including auto negotiation, duplex setting, ESRP host attach or don’t-count, and so on. All the ports in a load-sharing group must also be of the same bandwidth class.

The following rules apply:

- One group can contain up to 8 ports.
- The ports in the group do not need to be contiguous.
- A load share group must use ports that are all of the same maximum bandwidth capability.
- When using load sharing with the ESRP host attach feature, configure all ports in the same load-sharing group as host attach ports. When using load sharing with the ESRP don’t count feature, configure all ports in the same load-sharing group as don’t count ports. For further information about the ESRP host attach feature, see the *ExtremeWare Software User Guide*.

To define a load-sharing group, you assign a group of ports to a single, logical port number. To enable or disable a load-sharing group, use the following commands:

```
enable sharing <port> grouping <portlist> {algorithm {port-based | address-based |
round-robin}}
disable sharing [<port>]
```

NOTE

Do not disable a port that is part of a load-sharing group. Disabling the port prevents it from forwarding traffic, but still allows the link to initialize. As a result, a partner switch does not receive a valid indication that the port is not in a forwarding state, and the partner switch will continue to forward packets.

Load-Sharing Example

The following example defines a load-sharing group that contains ports 9 through 12, and uses the first port in the group as the master logical port 9:

```
enable sharing 9 grouping 9-12
```

In this example, logical port 9 represents physical ports 9 through 12.

When using load sharing, you should always reference the master logical port of the load-sharing group (port 9 in the previous example) when configuring or viewing VLANs. VLANs configured to use other ports in the load-sharing group will have those ports deleted from the VLAN when load sharing becomes enabled.

Verifying the Load-Sharing Configuration

The screen output resulting from the show ports sharing command lists the ports that are involved in load sharing and the master logical port identity.

```
---------------------------------
 * Summit400-48t:46 # show ports sharing Load Sharing Monitor
Config    Current    Ld Share    Ld Share   Link     Link
Master    Master     Type        Group      Status   Ups
=========================================================
  37       37           a          37          A     1
                        a          38          R     0
                        a          39          A     1
                        a          40          A     1
                        a          41          A     1
                        a          42          A     1

Link Status: (A) Active, (D) Disabled, (LB) Loopback, (ND) Not Distributing
             (NP) Not Present, (R) Ready

Ld Share Type: (a) address based
```

Switch Port-Mirroring

Port-mirroring configures the switch to copy all traffic associated with one or more ports. The monitor port can be connected to a network analyzer or RMON probe for packet analysis. The system uses a traffic filter that copies a group of traffic to the monitor port.

NOTE

Port mirroring is not supported with CPU-generated traffic.

You can define the traffic filter based on the physical port. All data that traverses the port, regardless of VLAN configuration, is copied to the monitor port.

Up to eight mirroring filters and one monitor port can be configured. Once a port is specified as a monitor port, it cannot be used for any other function.

NOTE

Frames that contain errors are not mirrored.

The mirrored port only transmits tagged frames. This allows you to mirror multiple ports or VLANs to a mirror port, while preserving the ability of a single protocol analyzer to track and differentiate traffic within a broadcast domain (VLAN) and across broadcast domains (for example, across VLANs when routing).

To enable port mirroring, use the following command:

```
enable mirroring to port [<port>] tagged
```

To configure the switch for port mirroring, use the following command:

```
configure mirroring add [<mac_address> | vlan <vlan name> {ports <port number>} |
ports <portnumber> {vlan <vlan name>}]
```

NOTE

When a mirrored port is configured, the forwarding database for items being mirrored (e.g., ports or VLANs) is automatically cleared if the link status on the mirrored port changes. This clearing results in some temporary flooding until the normal learning process completes. Removing or inserting a probe device into the mirror port may appear to cause flooding, but this temporary condition is normal.

Summit 400 Switch Port-Mirroring Example

The following example selects port 3 as the mirror port and sends all traffic coming into or out of the switch on port 1 to the mirror port:

```
enable mirroring to port 3 tagged
configure mirroring add port 1
```

The following example sends all traffic coming into or out of the switch on port 1 of *default* to the mirror port:

```
configure mirroring add port 1 default
```

Extreme Discovery Protocol

The Extreme Discovery Protocol (EDP) is used to gather information about neighbor Extreme Networks switches. EDP is used by the switches to exchange topology information. EDP is also used by the Extreme Standby Router Protocol (ESRP). Information communicated using EDP includes:

- Switch MAC address (switch ID).
- Switch software version information.
- Switch IP address.
- Switch VLAN-IP information.
- Switch port number.

EDP is enabled on all ports by default.

To disable EDP on one or more ports, use the following command:

```
disable edp ports [<portlist> | all]
```

To enable EDP on specified ports, use the following command:

```
enable edp ports [<portlist> | all]
```

To view EDP port information on the switch, use the following command:

```
show edp
```

Configuring Automatic Failover for Combination Ports

The Summit 400-48t allows you to configure all of the 1 Gigabit fiber ports and the first four 10/100/1000 copper ports for redundancy. For an introduction to this feature, see “Uplink Redundancy” on page 27.

The selection of whether a copper or fiber connection is determined by the order in which the cables are first inserted into the switch. For example, if you inserted a SFP connector into 1X and then a Ethernet cable into port 1, the fiber port becomes the primary uplink port and port 1 becomes the redundant port.

Hardware determines when a link is lost and swaps the primary and redundant ports to maintain stability. After a failover occurs, the switch keeps or sticks with the current port assignment until there is another failure or a user changes the assignment using the CLI. To change the uplink failover assignment, use the following command:

```
configure ports <nnn> preferred-medium {copper} | {fiber} |[force]
```

Using the `force` option disables automatic failover. If you force the preferred-medium to fiber and the fiber link goes away, the copper link is not used, even if available.

Automatic Failover Examples

If we can establish port 4 as the primary uplink and port 4X as the redundant uplink port using the CLI:

```
configure ports 4 preferred-medium copper
```

Port 4 becomes the primary uplink until a failure occurs on that link. At that time, fiber port 4X becomes the primary uplink and port 4 becomes the redundant port. This assignment stays in place until the next failure. However, if the 4X port is currently the primary medium when the command is issued, the command does not have an immediate effect.

In the next example, we force the switch to immediately start using the fiber port (if it currently has a link):

```
configure ports 3 preferred-medium fiber force
```

In this example, port 3X becomes the only uplink port. If the preferred-medium was copper when the command is issued, the switch immediately switches over and begins using the fiber port. If a failure occurs on the fiber port, the switch does not use the copper port as a redundant link. To allow the redundant uplink feature to be used again, issue this command:

```
configure ports 3 preferred-medium fiber
```

5 Virtual LANs (VLANs)

This chapter covers the following topics:

Setting up Virtual Local Area Networks (VLANs) on the switch eases many time-consuming tasks of network administration while increasing efficiency in network operations.

Overview of Virtual LANs

The term “VLAN” is used to refer to a collection of devices that communicate as if they were on the same physical LAN. Any set of ports (including all ports on the switch) is considered a VLAN. LAN segments are not restricted by the hardware that physically connects them. The segments are defined by flexible user groups you create with the command line interface.

Benefits

Implementing VLANs on your networks has the following advantages:

- **VLANs help to control traffic—**With traditional networks, congestion can be caused by broadcast traffic that is directed to all network devices, regardless of whether they require it. VLANs increase the efficiency of your network because each VLAN can be set up to contain only those devices that must communicate with each other.
- **VLANs provide extra security—**Devices within each VLAN can only communicate with member devices in the same VLAN. If a device in VLAN *Marketing* must communicate with devices in VLAN *Sales*, the traffic must cross a routing device.
- **VLANs ease the change and movement of devices—**With traditional networks, network administrators spend much of their time dealing with moves and changes. If users move to a different subnetwork, the addresses of each endstation must be updated manually.

Types of VLANs

VLANs can be created according to the following criteria:

- Physical port
- 802.1Q tag
- Ethernet, LLC SAP, or LLC/SNAP Ethernet protocol type
- MAC address
- A combination of these criteria

Port-Based VLANs

In a port-based VLAN, a VLAN name is given to a group of one or more ports on the switch. All ports are members of the port-based VLAN *default*. Before you can add any port to another port-based VLAN, you must remove it from the default VLAN, unless the new VLAN uses a protocol other than the default protocol *any*. A port can be a member of only one port-based VLAN.

On the Summit 400 switch in Figure 11, ports 9 through 14 are part of VLAN *Marketing*; ports 25 through 29 are part of VLAN *Sales*; and ports 21 through 24 and 30 through 32 are in VLAN *Finance*.

Figure 11: Example of a port-based VLAN on the Summit 400 switch

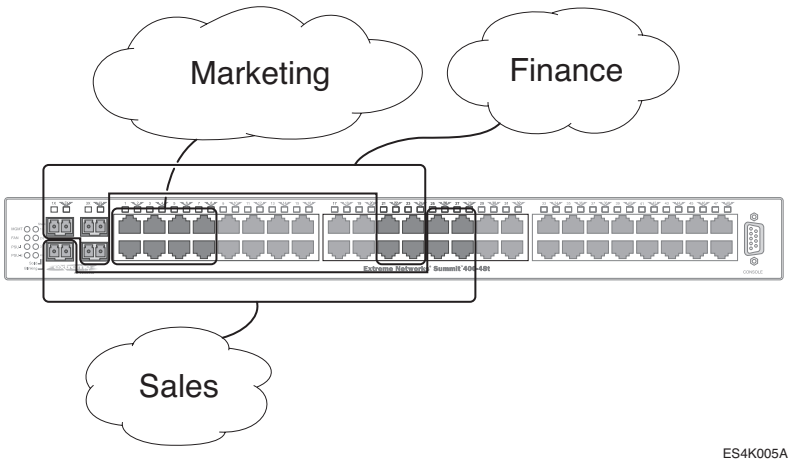

For the members of the different IP VLANs to communicate, the traffic must be routed by the switch. This means that each VLAN must be configured as a router interface with a unique IP address.

Spanning Switches with Port-Based VLANs

To create a port-based VLAN that spans two switches, you must do two things:

1 Assign the port on each switch to the VLAN.

2 Cable the two switches together using one port on each switch per VLAN.

Figure 12 illustrates a single VLAN that spans a BlackDiamond switch and a Summit 400 switch. All ports on the BlackDiamond switch belong to VLAN *Sales*. Port 1X and ports 1 through 28 on the Summit 400 switch also belong to VLAN *Sales*. The two switches are connected using slot 8, port 4 on system 1 (the BlackDiamond switch), and port 1X on system 2 (the Summit 400 switch).

Figure 12: Single port-based VLAN spanning two switches

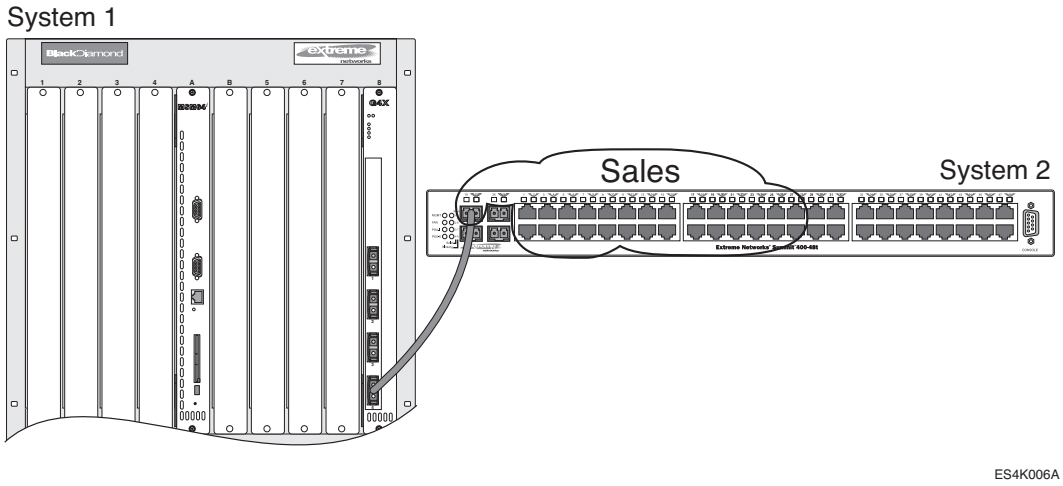

To create multiple VLANs that span two switches in a port-based VLAN, a port on system 1 must be cabled to a port on system 2 for each VLAN you want to have span across the switches. At least one port on each switch must be a member of the corresponding VLANs, as well.

Figure 13 illustrates two VLANs spanning two switches. On system 2, port 1X and ports 25 through 28 are part of VLAN *Accounting*; ports 21 through 24 and ports 2X through 4X are part of VLAN *Engineering*. On system 1, all ports on slot 1 are part of VLAN *Accounting*; all ports on slot 8 are part of VLAN *Engineering*.

Figure 13: Two port-based VLANs spanning two switches

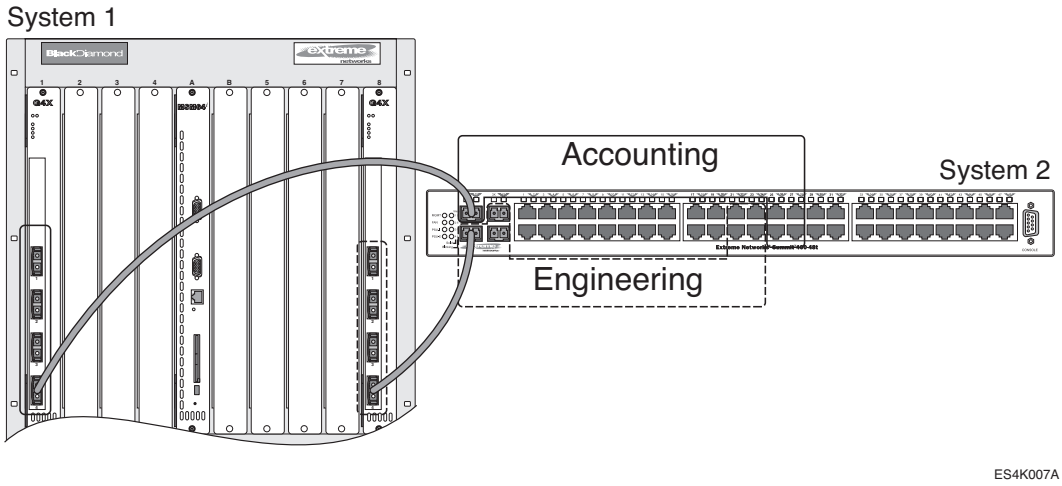

VLAN *Accounting* spans system 1 and system 2 by way of a connection between system 2, port 1X and system 1, slot 1, port 6. VLAN *Engineering* spans system 1 and system 2 by way of a connection between system 2, port 2X, and system 1, slot 8, port 6.

Using this configuration, you can create multiple VLANs that span multiple switches, in a daisy-chained fashion. Each switch must have a dedicated port for each VLAN. Each dedicated port must be connected to a port that is a member of its VLAN on the next switch.

Tagged VLANs

Tagging is a process that inserts a marker (called a *tag*) into the Ethernet frame. The tag contains the identification number of a specific VLAN, called the *VLANid*.

NOTE

The use of 802.1Q tagged packets may lead to the appearance of packets slightly bigger than the current IEEE 802.3/Ethernet maximum of 1,518 bytes. This may affect packet error counters in other devices, and may also lead to connectivity problems if non-802.1Q bridges or routers are placed in the path.

Uses of Tagged VLANs

Tagging is most commonly used to create VLANs that span switches. The switch-to-switch connections are typically called *trunks*. Using tags, multiple VLANs can span multiple switches using one or more trunks. In a port-based VLAN, each VLAN requires its own pair of trunk ports, as shown in Figure 13. Using tags, multiple VLANs can span two switches with a single trunk.

Another benefit of tagged VLANs is the ability to have a port be a member of multiple VLANs. This is particularly useful if you have a device (such as a server) that must belong to multiple VLANs. The device must have a NIC that supports 802.1Q tagging.

A single port can be a member of only one port-based VLAN. All additional VLAN membership for the port must be accompanied by tags. In addition to configuring the VLAN tag for the port, the server must have a *Network Interface Card (NIC)* that supports 802.1Q tagging.

Assigning a VLAN Tag

Each VLAN may be assigned an 802.1Q VLAN tag. As ports are added to a VLAN with an 802.1Q tag defined, you decide whether each port will use tagging for that VLAN. The default mode of the switch is to have all ports assigned to the VLAN named *default* with an 802.1Q VLAN tag (VLANid) of 1 assigned.

Not all ports in the VLAN must be tagged. As traffic from a port is forwarded out of the switch, the switch determines (in real time) if each destination port should use tagged or untagged packet formats for that VLAN. The switch adds and strips tags, as required, by the port configuration for that VLAN.

NOTE

Packets arriving tagged with a VLANid that is not configured on a port will be discarded.

Figure 14 illustrates the physical view of a network that uses tagged and untagged traffic.

Figure 14: Physical diagram of tagged and untagged traffic

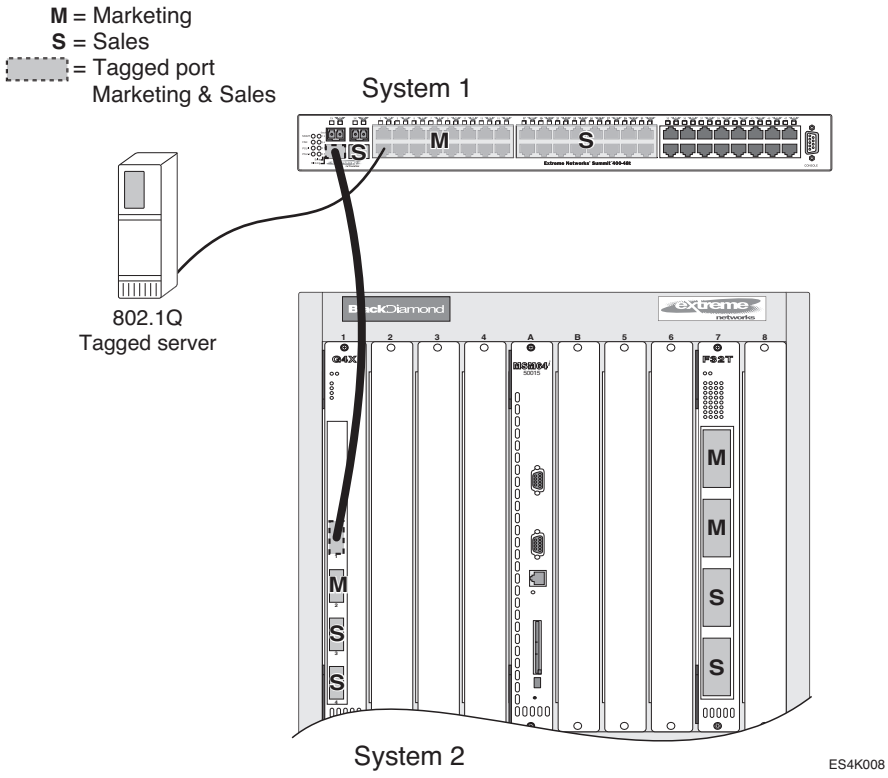

Figure 15 is a logical diagram of the same network.

Figure 15: Logical diagram of tagged and untagged traffic

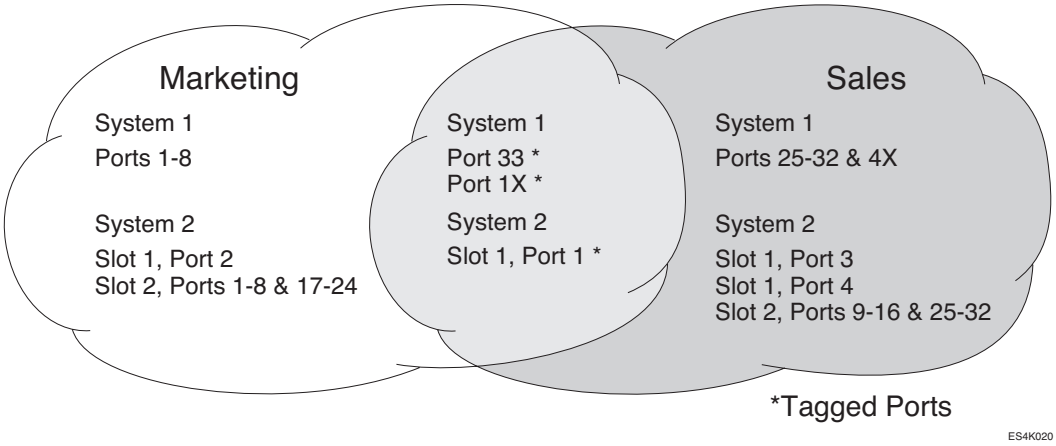

In Figure 14 and Figure 15:

- The trunk port on each switch carries traffic for both VLAN *Marketing* and VLAN *Sales.*
- The trunk port on each switch is tagged.
- The server connected to port 33 on system 1 has a NIC that supports 802.1Q tagging.
- The server connected to port 33 on system 1 is a member of both VLAN *Marketing* and VLAN *Sales.*
- All other stations use untagged traffic.

As data passes out of the switch, the switch determines if the destination port requires the frames to be tagged or untagged. All traffic coming from and going to the server is tagged. Traffic coming from and going to the trunk ports is tagged. The traffic that comes from and goes to the other stations on this network is not tagged.

Mixing Port-Based and Tagged VLANs

You can configure the switch using a combination of port-based and tagged VLANs. A given port can be a member of multiple VLANs, with the stipulation that only one of its VLANs uses untagged traffic. In other words, a port can simultaneously be a member of one port-based VLAN and multiple tag-based VLANs.

NOTE

For the purposes of VLAN classification, packets arriving on a port with an 802.1Q tag containing a VLANid of zero are treated as untagged.

VLAN Names

Each VLAN is given a name that can be up to 32 characters. VLAN names can use standard alphanumeric characters. The following characters are not permitted in a VLAN name:

- Space
- Comma

- Quotation mark

VLAN names must begin with an alphabetical letter. Quotation marks can be used to enclose a VLAN name that includes special characters, including single quotation marks or commas. Spaces may not be included, even within quotation marks. For example, the names *test*, *test1*, and *test_15* are acceptable VLAN names. The names "*test&5*" and "*joe's*" may be used if enclosed in quotation marks. Names such as "*5test*" or "*test 5*" are not permitted.

VLAN names can be specified using the tab key for command completion.

VLAN names are locally significant. That is, VLAN names used on one switch are only meaningful to that switch. If another switch is connected to it, the VLAN names have no significance to the other switch.

NOTE

You should use VLAN names consistently across your entire network.

Default VLAN

The switch ships with one default VLAN that has the following properties:

- The VLAN name is *default.*
- It contains all the ports on a new or initialized switch.
- The default VLAN is untagged on all ports. It has an internal VLANid of 1.

Renaming a VLAN

To rename an existing VLAN, use the following command:

```
configure vlan <old_name> name <new_name>
```

The following rules apply to renaming VLANs:

- Once you change the name of the default VLAN, it cannot be changed back to *default*.
- You cannot create a new VLAN named *default*.
- You cannot change the VLAN name *MacVlanDiscover*. Although the switch accepts a name change, once it is rebooted, the original name is recreated.

Configuring VLANs on the Switch

This section describes the commands associated with setting up VLANs on the switch. Configuring a VLAN involves the following steps:

1 Create and name the VLAN.

2 Assign an IP address and mask (if applicable) to the VLAN, if needed.

NOTE

Each IP address and mask assigned to a VLAN must represent a unique IP subnet. You cannot configure the same IP subnet on different VLANs.

NOTE

If you plan to use this VLAN as a control VLAN for an EAPS domain, do NOT assign an IP address to the VLAN.

3 Assign a VLANid, if any ports in this VLAN will use a tag.

4 Assign one or more ports to the VLAN.

As you add each port to the VLAN, decide if the port will use an 802.1Q tag.

VLAN Configuration Examples

The following example creates a port-based VLAN named *accounting*, assigns the IP address 132.15.121.1, and assigns ports 1, 2, 3, and 6:

```
create vlan accounting
configure accounting ipaddress 132.15.121.1
configure default delete port 1-3,6
configure accounting add port 1-3,6
```

NOTE

Because VLAN names are unique, you do not need to enter the keyword `vlan` *after you have created the unique VLAN name. You can use the VLAN name alone.*

The following example creates a tag-based VLAN named *video*. It assigns the VLANid 1000. Ports 4 through 8 are added as tagged ports to the VLAN.

```
create vlan video
configure video tag 1000
configure video add port 4-8 tagged
```

The following example creates a VLAN named *sales*, with the VLANid 120. The VLAN uses both tagged and untagged ports. Ports 1 through 3 are tagged, and ports 4 and 7 are untagged. Note that when not explicitly specified, ports are added as untagged.

```
create vlan sales
configure sales tag 120
configure sales add port 1-3 tagged
configure default delete port 4,7
configure sales add port 4,7
```

Displaying VLAN Settings

To display VLAN settings, use the following command:

```
show vlan {<vlan name> | detail | stats {vlan} <vlan name>}
```

The `show` command displays summary information about a specific VLAN, which includes:

- Name.
- VLANid.

- How the VLAN was created.
- IP address.
- STPD information.
- Protocol information.
- QoS profile information.
- Ports assigned.
- Tagged/untagged status for each port.
- How the ports were added to the VLAN.
- Number of VLANs configured on the switch.

Use the `detail` option to display the detailed format.

MAC-Based VLANs

MAC-Based VLANs allow physical ports to be mapped to a VLAN based on the source MAC address learned in the FDB. This feature allows you to designate a set of ports that have their VLAN membership dynamically determined by the MAC address of the end station that plugs into the physical port. You can configure the source MAC address-to-VLAN mapping either offline or dynamically on the switch. For example, you could use this application for a roaming user who wants to connect to a network from a conference room. In each room, the user plugs into one of the designated ports on the switch and is mapped to the appropriate VLAN. Connectivity is maintained to the network with all of the benefits of the configured VLAN in terms of QoS, routing, and protocol support.

MAC-Based VLAN Guidelines

When using the MAC-to-VLAN mapping, consider the following guidelines:

- A port can only accept connections from an endstation/host and should not be connected to a layer-2 repeater device. Connecting to a layer-2 repeater device can cause certain addresses to not be mapped to their respective VLAN if they are not correctly configured in the MAC-VLAN configuration database. If a repeater device is connected to a MAC-Based VLAN port, and the configured MAC-to-VLAN mapped station enters on the repeater, any endstation that is attached to the repeater can be mapped to that VLAN while the configured endstation is active in that VLAN. Upon removal of the configured MAC-to-VLAN endstation, all other endstations lose connectivity.
- Groups are used as a security measure to allow a MAC address to enter into a VLAN only when the group mapping matches the port mapping.

  As an example, the following configuration allows MAC 00:00:00:00:00:aa to enter into the VLAN only on ports 10 and 11 because of membership in group 100:

```
* Summit400 # show mac-vlan
Port    Vlan                Group         State
10      MacVlanDiscover     100           Discover
11      MacVlanDiscover     100           Discover
12      MacVlanDiscover     any           Discover
13      MacVlanDiscover     any           Discover
14      MacVlanDiscover     any           Discover
Total Entries in Database:2
 Mac                Vlan     Group
00:00:00:00:00:aa   sales    100
```

```
00:00:00:00:00:01    sales    any
2 matching entries
```

- The group “any” is equivalent to the group “0”. Ports that are configured as “any” allow any MAC address to be assigned to a VLAN, regardless of group association.
- Partial configurations of the MAC to VLAN database can be downloaded to the switch using the timed download configuration feature.

MAC-Based VLAN Limitations

The following list contains the limitations of MAC-based VLANs:

- Ports participating in MAC VLANs must first be removed from any static VLANs.
- The MAC- to-VLAN mapping can only be associated with VLANs that exist on the switch.
- A MAC address cannot be configured to associate with more than 1 VLAN. If this is attempted, the MAC address is associated with the most recent VLAN entry in the MAC-to-VLAN database.
- The feature is intended to support one client per physical port. Once a client MAC address has successfully registered, the VLAN association remains until the port connection is dropped or the FDB entry ages out.

MAC-Based VLAN Example

In this following example, three VLANs are created: *engineering*, *marketing*, and *sales*. A single MAC address is associated with each VLAN. The MAC address 00:00:00:00:00:02 has a group number of “any” or “0” associated with it, allowing it to be plugged into any port that is in MacVlanDiscover mode (ports 10-15 in this case). The MAC address 00:00:00:00:00:01 has a group number of 10 associated with it, and can only be assigned to a VLAN if inserted into ports 16 or 17. The MAC address 00:00:00:00:00:03 has a group number of 200 associated with it and can only be inserted into ports 18 through 20.

```
enable mac-vlan mac-group any ports 10-15
enable mac-vlan mac-group 10 ports 16-17
enable mac-vlan mac-group 200 ports 18-20
configure mac-vlan add mac-address 00:00:00:00:00:01 mac-group 10 engineering
configure mac-vlan add mac-address 00:00:00:00:00:02 mac-group any marketing
configure mac-vlan add mac-address 00:00:00:00:00:03 mac-group 200 sales
```

Timed Configuration Download for MAC-Based VLANs

To allow centralized control of MAC-based VLANs over multiple switches, a timed TFTP configuration download allows you to download incremental configuration files from a primary or secondary server at specified time intervals. The timed downloads are configurable in 24 hour intervals. When a switch reboots, the configuration is automatically downloaded immediately after booting, per the configured primary and secondary servers.

To configure the primary and/or secondary server and file name, use the following command:

```
configure download server [primary | secondary] [<ip address> | <hostname>] <filename>
```

To enable timed interval downloads, use the following command:

```
download configuration every <time>
```

To display timed download information, use the following command:

```
show switch
```

Example

In relation to MAC-based VLANs, the downloaded file is an ASCII file that consists of CLI commands used to configure the most recent MAC-to-VLAN database. This feature is different from the normal download configuration command in that it allows incremental configuration without the automatic rebooting of the switch.

The following example shows an incremental configuration file for MAC-based VLAN information that updates the database and saves changes:

```
configure mac-vlan add mac-address 00:00:00:00:00:01 mac-group any engineering
configure mac-vlan add mac-address 00:00:00:00:ab:02 mac-group any engineering
configure mac-vlan add mac-address 00:00:00:00:cd:04 mac-group any sales
.
.
configure mac-vlan add mac-address 00:00:00:00:ab:50 mac-group any sales
configure mac-vlan add mac-address 00:00:00:00:cd:60 mac-group any sales
save
```

6 Forwarding Database (FDB)

This chapter describes the following topics:

Overview of the FDB

The switch maintains a database of all media access control (MAC) addresses received on all of its ports. It uses the information in this database to decide whether a frame should be forwarded or filtered.

FDB Contents

Each FDB entry consists of:

- The MAC address of the device
- An identifier for the port and VLAN on which it was received
- The age of the entry
- The number of IP FDB entries that use this MAC address as a next hop or last hop
- Flags

Frames destined for MAC addresses that are not in the FDB are flooded to all members of the VLAN.

How FDB Entries Get Added

Entries are added into the FDB in the following ways:

- The switch can learn entries by examining packets it receives. The system updates its FDB with the source MAC address from a packet, the VLAN, and the port identifier on which the source packet is received.

  The ability to learn MAC addresses can be enabled or disabled on a port-by-port basis. You can also limit the number of addresses that can be learned, or you can “lock down” the current entries and prevent additional MAC address learning.

- You can enter and update entries using the command line interface (CLI).
- Certain static entries are added by the system upon switch boot up.

FDB Entry Types

FDB entries may be dynamic or static, and may be permanent or non-permanent. The following describes the types of entries that can exist in the FDB:

- **Dynamic entries**—A dynamic entry is learned by the switch by examining packets to determine the source MAC address, VLAN, and port information. The switch then creates or updates an FDB entry for that MAC address. Initially, all entries in the database are dynamic, except for certain entries created by the switch at boot up.

  Dynamic entries are flushed and relearned (updated) when any of the following take place:

  — A VLAN is deleted.

  — A VLAN identifier (VLANid) is changed.

  — A port mode is changed (tagged / untagged).

  — A port is deleted from a VLAN.

  — A port is disabled.

  — A port enters blocking state.

  — A port QoS setting is changed.

  — A port goes down (link down).

  A *non-permanent dynamic entry* is initially created when the switch identifies a new source MAC address that does not yet have an entry in the FDB. The entry may then be updated as the switch continues to encounter the address in the packets it examines. These entries are identified by the “d” flag in `show fdb` output.

  A *permanent dynamic entry* is created by command through the CLI, but may then be updated as the switch encounters the MAC address in the packets that it examines. A permanent dynamic entry is typically used to associate QoS profiles with the FDB entry. Permanent dynamic entries are identified by the “p” and “d” flags in `show fdb` output.

  Both types of dynamic entries age—a dynamic entry will be removed from the FDB (aged-out) if the device does not transmit for a specified period of time (the aging time). This prevents the database from becoming full with obsolete entries by ensuring that when a device is removed from the network, its entry is deleted from the database. The aging time is configurable. For more information about setting the aging time, see “Configuring the FDB Aging Time” on page 103 later in this chapter.

- **Static entries**—A static entry does not age, and does not get updated through the learning process. It is maintained exactly as it was created. Conditions that cause dynamic entries to be updated, such as VLAN or port configuration changes, do not affect static entries.

  If the same MAC address is detected on another virtual port that is not defined in the static FDB entry for the MAC address, it is handled as a blackhole entry.

  A permanent static entry is created through the command line interface, and can be used to associate QoS profiles with a non-aging FDB entry. Permanent static entries are identified by the “s” and “p” flags in `show fdb` output.

  A *locked static entry* is an entry that was originally learned dynamically, but has been made static (locked) using the MAC address lock-down feature. It is identified by the “s” and “l” flags in `show fdb` output. See “Network Login” on page 146 for more information about MAC address lock-down.

Non-permanent static entries are created by the switch software for various reasons, typically upon switch boot up. They are identified by the "s" flag in `show fdb` output.

If the FDB entry aging time is set to zero, all entries in the database are considered static, non-aging entries. This means that they do not age, but they are still deleted if the switch is reset.

- **Permanent entries**—Permanent entries are retained in the database if the switch is reset or a power off/on cycle occurs. Permanent entries must be created by the system administrator through the command line interface. A permanent entry can either be a unicast or multicast MAC address.

  Permanent entries may be static, meaning they do not age or get updated, or they may be dynamic, meaning that they do age and can be updated via learning.

  Permanent entries can have QoS profiles associated with the MAC address. A different QoS profiles may be associated with the MAC address when it is a destination address (an egress QoS profile) than when it is a source address (ingress QoS profile).

  The Summit 400 can support a maximum of 64 permanent entries.

- **Blackhole entrie**s—A blackhole entry configures the switch to discard packets with a specified MAC address. Blackhole entries are useful as a security measure or in special circumstances where a specific source or destination address must be discarded. Blackhole entries may be created through the CLI, or they may be created by the switch when a port's learning limit has been exceeded.

  Blackhole entries are treated like permanent entries in the event of a switch reset or power off/on cycle. Blackhole entries are never aged out of the database.

Disabling MAC Address Learning

By default, MAC address learning is enabled on all ports. You can disable learning on specified ports using the following command:

```
disable learning ports <portlist>
```

If MAC address learning is disabled, only broadcast traffic, EDP traffic, and packets destined to a permanent MAC address matching that port number, are forwarded. Use this command in a secure environment where access is granted via permanent forwarding databases (FDBs) per port.

Associating QoS Profiles with an FDB Entry

You can associate QoS profiles with a MAC address (and VLAN) of a device by creating a permanent FDB entry and specifying QoS profiles for ingress or egress, or both. The permanent FDB entry can be either dynamic (it is learned and can be aged out) or static.

To associate a QoS profile with a dynamic FDB entry, use the following command:

```
create fdbentry [<mac_address> | any-mac] vlan <vlan name> dynamic ingress-qosprofile
<qosprofile>{ingress-qosprofile <inqosprofile>}]
```

This command associates QoS profiles with packets received from or destined for the specified MAC address, while still allowing the FDB entry to be dynamically learned. If you specify only the ingress QoS profile, the egress QoS profile defaults to none, and vice-versa. If both profiles are specified, the source MAC address of an ingress packet and the destination MAC address of an egress packet are examined for QoS profile assignment.

The FDB entry is not actually created until the MAC address is encountered as the source MAC address in a packet. Thus, initially the entry may not appear in the `show fdb` output. Once the entry has been

learned, it is created as a permanent dynamic entry, designated by “dpm” in the flags field of the `show fdb` output.

You can display permanent FDB entries, including their QoS profile associations by using the `permanent` option in the following command:

```
show fdb {<mac_address> | permanent | ports <portlist> | vlan <vlan name>}
```

To associate a QoS profile with a permanent FDB entry, use the following command:

```
create fdbentry <mac_address> vlan <vlan name> ports [<portlist> | all] {qosprofile
<qosprofile>}{ingress-qosprofile <inqosprofile>}
```

This entry will not be aged out, and no learning will occur. If the same MAC address is encountered through a virtual port not specified in the portlist, it will be handled as a blackhole entry.

Using the `any-mac` keyword, you can enable traffic from a QoS VLAN to have higher priority than 802.1p traffic. Normally, an 802.1p packet has a higher priority over the VLAN classification. In order to use this feature, you must create a wildcard permanent FDB entry named `any-mac` and apply the QoS profile to the individual MAC entry.

NOTE

For more information on QoS profiles, see Chapter 7.

FDB Configuration Examples

The following example adds a permanent static entry to the FDB:

```
create fdbentry 00:E0:2B:12:34:56 vlan marketing port 3:4
```

The permanent entry has the following characteristics:

- MAC address is 00:E0:2B:12:34:56.
- VLAN name is *marketing*.
- Slot number for this device is 3.
- Port number for this device is 4.

If the MAC address 00:E0:2B:12:34:56 is encountered on any port/VLAN other than VLAN *marketing*, port 3:4, it will be handled as a blackhole entry, and packets from that source will be dropped.

This example associates the QoS profile *qp2* with a dynamic entry for the device at MAC address 00:A0:23:12:34:56 on VLAN *net34* that will be learned by the FDB:

```
create fdbentry 00:A0:23:12:34:56 vlan net34 dynamic qosprofile qp2
```

This entry has the following characteristics:

- MAC address is 00:A0:23:12:34:56.
- VLAN name is *net34*.
- The entry will be learned dynamically.
- QoS profile *qp2* will be applied as an egress QoS profile when the entry is learned.

Overriding 802.1p Priority

This example associates the QoS profile *qp5* with the wildcard permanent FDB entry *any-mac* on VLAN v110:

```
create fdbentry any-mac vlan v110 dynamic ingress-qosprofile qp5
```

Configuring the FDB Aging Time

You can configure the again time for dynamic FDB entries using the following command:

```
configure fdb agingtime <seconds>
```

If the aging time is set to zero, all aging entries in the database are defined as static, nonaging entries. This means they will not age out, but non-permanent static entries can be deleted if the switch is reset.

Displaying FDB Entries

To display FDB entries, use the following command:

```
show fdb {<mac_address> | permanent | ports <portlist> | vlan <vlan name>}
```

where the following is true:

- `mac_address`—Displays the entry for a particular MAC address.
- `broadcast-mac`—Specifies the broadcast MAC address. May be used as an alternate to the colon-separated byte form of the address ff:ff:ff:ff:ff:ff
- `permanent`—Displays all permanent entries, including the ingress and egress QoS profiles.
- `ports <portlist>`—Displays the entries for a set of ports or slots and ports.
- `vlan <vlan name>`—Displays the entries for a VLAN.

With no options, the command displays all FDB entries.

See the *ExtremeWare 7.2e Command Reference Guide* for details of the commands related to the FDB.

7 Quality of Service (QoS)

This chapter covers the following topics:

Policy-based Quality of Service (QoS) is a feature of ExtremeWare and the Extreme switch architecture that allows you to specify different service levels for traffic traversing the switch. Policy-based QoS is an effective control mechanism for networks that have heterogeneous traffic patterns. Using Policy-based QoS, you can specify the service level that a particular traffic type receives.

Overview of Policy-Based Quality of Service

Policy-based QoS allows you to protect bandwidth for important categories of applications or specifically limit the bandwidth associated with less critical traffic. For example, if voice–over-IP traffic requires a reserved amount of bandwidth to function properly, using policy-based QoS, you can reserve sufficient bandwidth critical to this type of application. Other applications deemed less critical can be limited so as to not consume excessive bandwidth. The switch contains separate hardware queues on every physical port. The prioritization parameters that modify the forwarding behavior of the switch affect how the switch transmits traffic for a given hardware queue on a physical port. Up to eight physical queues per port are available.

NOTE

Policy-based QoS has no impact on switch performance. Using even the most complex traffic groupings has no cost in terms of switch performance.

Applications and Types of QoS

Different applications have different QoS requirements. The following applications are ones that you will most commonly encounter and need to prioritize:

- Voice applications
- Video applications
- Critical database applications
- Web browsing applications
- File server applications

General guidelines for each traffic type are given below and summarized in Table 19. Consider them as general guidelines and not strict recommendations. Once QoS parameters are set, you can monitor the performance of the application to determine if the actual behavior of the applications matches your expectations. It is very important to understand the needs and behavior of the particular applications you wish to protect or limit. Behavioral aspects to consider include bandwidth needs, sensitivity to latency and jitter, and sensitivity and impact of packet loss.

Voice Applications

Voice applications typically demand small amounts of bandwidth. However, the bandwidth must be constant and predictable because voice applications are typically sensitive to latency (inter-packet delay) and jitter (variation in inter-packet delay). The most important QoS parameter to establish for voice applications is minimum bandwidth, followed by priority.

Video Applications

Video applications are similar in needs to voice applications, with the exception that bandwidth requirements are somewhat larger, depending on the encoding. It is important to understand the behavior of the video application being used. For example, in the playback of stored video streams, some applications can transmit large amounts of data for multiple streams in one “spike,” with the expectation that the end-stations will buffer significant amounts of video-stream data. This can present a problem to the network infrastructure, because it must be capable of buffering the transmitted spikes

where there are speed differences (for example, going from Gigabit Ethernet to Fast Ethernet). Key QoS parameters for video applications include minimum bandwidth, priority, and possibly buffering (depending upon the behavior of the application).

Critical Database Applications

Database applications, such as those associated with ERP, typically do not demand significant bandwidth and are tolerant of delay. You can establish a minimum bandwidth using a priority less than that of delay-sensitive applications.

Web Browsing Applications

QoS needs for Web browsing applications cannot be generalized into a single category. For example, ERP applications that use a browser front-end may be more important than retrieving daily news information. Traffic groupings can typically be distinguished from each other by their server source and destinations. Most browser-based applications are distinguished by the dataflow being asymmetric (small dataflows from the browser client, large dataflows from the server to the browser client).

An exception to this may be created by some Java™ -based applications. In addition, Web-based applications are generally tolerant of latency, jitter, and some packet loss, however small packet-loss may have a large impact on perceived performance due to the nature of TCP. The relevant parameter for protecting browser applications is minimum bandwidth. The relevant parameter for preventing non-critical browser applications from overwhelming the network is maximum bandwidth.

File Server Applications

With some dependencies on the network operating system, file serving typically poses the greatest demand on bandwidth, although file server applications are very tolerant of latency, jitter, and some packet loss, depending on the network operating system and the use of TCP or UDP.

NOTE

Full-duplex links should be used when deploying policy-based QoS. Half-duplex operation on links can make delivery of guaranteed minimum bandwidth impossible.

Table 19 summarizes QoS guidelines for the different types of network traffic.

Table 19: Traffic Type and QoS Guidelines

Traffic Type	Key QoS Parameters
Voice	Minimum bandwidth, priority
Video	Minimum bandwidth, priority, buffering (varies)
Database	Minimum bandwidth
Web browsing	Minimum bandwidth for critical applications, maximum bandwidth for non-critical applications
File server	Minimum bandwidth

Configuring QoS

To configure QoS, you define how your switch responds to different categories of traffic by creating and configuring QoS profiles. You then group traffic into categories (according to application, as previously discussed) and assign each category to a QoS profile. Configuring QoS is a three-step process:

1 Configure the QoS profile.

QoS profile—A class of service that is defined through prioritization settings. The level of service that a particular type of traffic or traffic grouping receives is determined by assigning it to a QoS profile.

2 Create traffic groupings.

Traffic grouping—A classification or traffic type that has one or more attributes in common, such as a physical port. You assign traffic groupings to QoS profiles to modify switch forwarding behavior. Traffic groupings transmitting out the same port that are assigned to a particular QoS profile share the assigned prioritization characteristics, and hence share the class of service.

3 Monitor the performance of the application with the QoS monitor to determine whether the policies are meeting the desired results.

The next sections describe each of these QoS components in detail.

QoS Profiles

A QoS profile defines a class of service by specifying traffic behavior attributes, such as bandwidth. The parameters that make up a QoS profile include:

- **Priority**—The level of priority assigned to a hardware queue on a physical port. There are eight different available priority settings. By default, each of the default QoS profiles is assigned a unique priority. You would use prioritization when two or more hardware queues on the same physical port are contending for transmission on the same physical port, only after their respective bandwidth management parameters have been satisfied.
  - — When configured to do so, the priority of a QoS profile can determine the 802.1p bits used in the priority field of a transmitted packet (described later).
  - — The priority of a QoS profile determines the DiffServ code point value used in an IP packet when the packet is transmitted (described later).

A QoS profile does not alter the behavior of the switch until it is assigned to a traffic grouping. Recall that QoS profiles are linked to hardware queues. There are multiple hardware queues per physical port. By default, a QoS profile links to the identical hardware queue across all the physical ports of the switch.

The default QoS profiles cannot be deleted. Administrators do not have the authority to create a QoS profile. Also by default, a QoS profile maps directly to a specific hardware queue across all physical ports. The settings for the default QoS parameters are summarized in Table 20.

Table 20: QoS Parameters

Profile Name	Hardware Queue	Priority	Buffer	Minimum Bandwidth	Maximum Bandwidth
Qp1	Q0	Low	0	0%	100%

Table 20: QoS Parameters (Continued)

Qp2	Q1	Lowhi	0	0%	100%
Qp3	Q2	Normal	0	0%	100%
Qp4	Q3	Normalhi	0	0%	100%
Qp5	Q4	Medium	0	0%	100%
Qp6	Q5	Mediumhi	0	0%	100%
Qp7	Q6	High	0	0%	100%
Qp8	Q7	Highhi	0	0%	100%

Traffic Groupings

A *traffic grouping* is a classification of traffic that has one or more attributes in common. Traffic is typically grouped based on the applications discussed starting on page 106.

Traffic groupings are separated into the following categories for discussion:

- IP-based information, such as IP source/destination and TCP/UDP port information
- Destination MAC (MAC QoS groupings)
- Explicit packet class of service information, such as 802.1p or DiffServ (IP TOS)
- Physical/logical configuration (physical source port or VLAN association)

In the event that a given packet matches two or more grouping criteria, there is a predetermined precedence for which traffic grouping will apply. In general, the more specific traffic grouping takes precedence. By default, all traffic groupings are placed in the QoS profile Qp1. The supported traffic groupings are listed in Table 21. The groupings are listed in order of precedence (highest to lowest). The four types of traffic groupings are described in detail on the following pages.

Table 21: Traffic Groupings by Precedence

IP Information (Access Lists) Groupings
• Access list precedence determined by user configuration
Destination Address MAC-Based Groupings
• Permanent • Dynamic • Blackhole
Explicit Packet Class of Service Groupings
• DiffServ (IP TOS) • 802.1P
Physical/Logical Groupings
• VLAN • Source port

IP-Based Traffic Groupings

IP-based traffic groupings are based on any combination of the following items:

- IP source or destination address
- TCP or UDP protocols
- TCP/UDP port information

IP-based traffic groupings are defined using access lists. Access lists are discussed in detail in “IP Access Lists (ACLs)” on page 138. By supplying a named QoS profile at the end of the access list command syntax, you can prescribe the bandwidth management and priority handling for that traffic grouping. This level of packet filtering has no impact on performance.

For example, to create an IP-based traffic grouping, use the following commands:

```
create access-mask amask source-ip / 24 dest-ip / 24 precedence 2000

create access-list alist "amask" dest-ip 10.1.2.1/24 source-ip 10.1.1.1/24 permit
qosprofile qp3
```

To create a MAC-based traffic grouping, use this command:

```
create fdbentry 00 : 11 : 22 : 33 : 44 : 55 vlan "Default" dynamic qosprofile "QP3"
```

MAC-Based Traffic Groupings

QoS profiles can be assigned to destination MAC addresses. MAC-based traffic groupings are configured using the `create fdb...` command:

The MAC address options, defined below, are as follows:

- Permanent
- Dynamic
- Blackhole

NOTE

On the Summit 400 broadcast MAC entries may not be associated with a QoS.

Permanent MAC addresses

Permanent MAC addresses can be assigned a QoS profile whenever traffic is destined to the MAC address. This can be done when you create a permanent FDB entry using the following command:

```
create fdbentry <mac_address> vlan <vlan name> ports [<portlist> | all] {qosprofile
<qosprofile>}{ingress-qosprofile <inqosprofile>}
```

For example:

```
create fdbentry 00:11:22:33:44:55 vlan default port 4:1 qosprofile qp2
```

Dynamic MAC Addresses

Dynamic MAC addresses can be assigned a QoS profile whenever traffic is coming from or going to the MAC address. This is done using the following command:

```
create fdbentry [<mac_address> | any-mac] vlan <vlan name> dynamic ingress-qosprofile
<qosprofile>{ingress-qosprofile <inqosprofile>}
```

For any port on which the specified MAC address is learned in the specified VLAN, the port is assigned the specified QoS profile. For example:

```
create fdbentry 00 : 11 : 22 : 33 : 44 : 55 vlan "Default" dynamic ingress-qosprofile
"QP1" qosprofile qp2
```

The QoS profile is assigned when the MAC address is learned. If a client's location moves, the assigned QoS profile moves with the device. If the MAC address entry already exists in the FDB, you can clear the forwarding database so that the QoS profile can be applied when the entry is added again. Use the following command to clear the FDB:

```
clear fdb
```

Blackhole MAC Address

Using the `blackhole` option configures the switch to not forward any packets to the destination MAC address on any ports for the VLAN specified. The `blackhole` option is configured using the following command:

```
create fdbentry <mac_address> vlan <vlan name> blackhole {source-mac | dest-mac |
both}
```

For example:

```
create fdbentry 00:11:22:33:44:55 vlan default blackhole
```

Verifying MAC-Based QoS Settings

To verify any of the MAC-based QoS settings, use either the command

```
show fdb permanent
```

or the command

```
show qosprofile {<qosprofile>} {port <portlist>}
```

Explicit Class of Service (802.1p and DiffServ) Traffic Groupings

This category of traffic groupings describes what is sometimes referred to as *explicit packet marking*, and refers to information contained within a packet intended to explicitly determine a class of service. That information includes:

- IP DiffServ code points, formerly known as IP TOS bits
- Prioritization bits used in IEEE 802.1p packets

An advantage of explicit packet marking is that the class of service information can be carried throughout the network infrastructure, without repeating what can be complex traffic grouping policies at each switch location. Another advantage is that end stations can perform their own packet marking

on an application-specific basis. Extreme switch products have the capability of observing and manipulating packet marking information with no performance penalty.

The documented capabilities for 802.1p priority markings or DiffServ capabilities (if supported) are not impacted by the switching or routing configuration of the switch. For example, 802.1p information can be preserved across a routed switch boundary and DiffServ code points can be observed or overwritten across a layer 2 switch boundary.

Configuring 802.1p Priority

Extreme switches support the standard 802.1p priority bits that are part of a tagged Ethernet packet. The 802.1p bits can be used to prioritize the packet, and assign it to a particular QoS profile.

When a packet arrives at the switch, the switch examines the 802.1p priority field maps it to a specific hardware queue when subsequently transmitting the packet. The 802.1p priority field is located directly following the 802.1Q type field, and preceding the 802.1Q VLAN ID, as shown in Figure 16.

Figure 16: Ethernet packet encapsulation

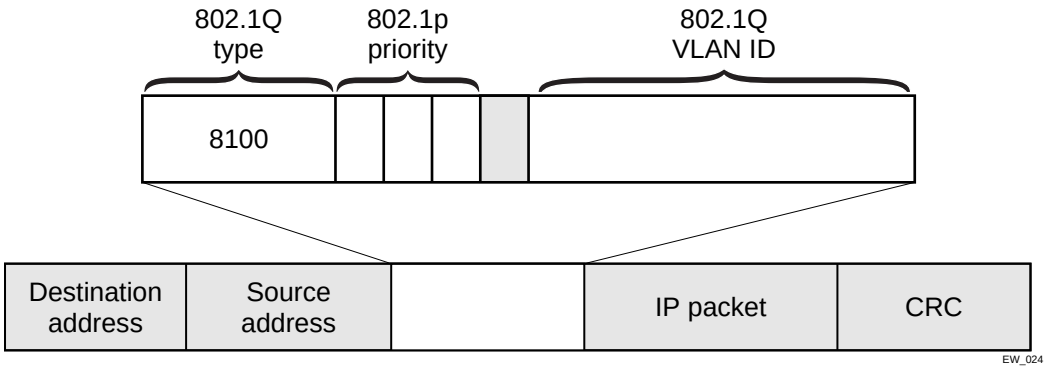

Observing 802.1p Information

When ingress traffic that contains 802.1p prioritization information is detected by the switch, the traffic is mapped to various hardware queues on the egress port of the switch. Eight hardware queues are supported. The transmitting hardware queue determines the priority characteristics used when transmitting packets.

To control the mapping of 802.1p prioritization values to hardware queues, 802.1p prioritization values can be mapped to a QoS profile. The default mapping of each 802.1p priority value to QoS profile is shown in Table 22.

Table 22: 802.1p Priority Value-to-QoS Profile Default Mapping

Priority Value	QoS Profile
0	Qp1
1	Qp2
2	Qp3
3	Qp4
4	Qp5
5	Qp6

Table 22: 802.1p Priority Value-to-QoS Profile Default Mapping (Continued)

Priority Value	QoS Profile
6	Qp7
7	Qp8

Configuring 802.1p Priority For Slow Path Traffic

Some traffic can originate on the switch, for example Ping or Telnet packets. This traffic comes from the switch CPU and is referred to as slow path traffic. This traffic is internally tagged with an 802.1p priority of 7, by default, and egresses the VLAN through the highest queue. If you want to set a different tag (and priority) use the following command to set the priority to a number between 0 and 7:

```
configure vlan <vlan name> priority <priority>
```

Other traffic transported across the switch and VLAN will not be changed, in other words, the 802.1p values will not be affected by the VLAN priority setting.

Replacing 802.1p Priority Information

By default, 802.1p priority information is not replaced or manipulated, and the information observed on ingress is preserved when transmitting the packet. This behavior is not affected by the switching or routing configuration of the switch.

However, the switch is capable of replacing the 802.1p priority information. To replace 802.1p priority information, you will use an access list to set the 802.1p value. See “IP Access Lists (ACLs)” on page 138, for more information on using access lists. You will use the `set dot1p <dot1p_value>` parameter of the `create access-list` command to replace the value. The packet is then placed on the queue that corresponds to the new 802.1p value.

Configuring DiffServ

Contained in the header of every IP packet is a field for IP Type of Service (TOS), now also called the DiffServ field. The TOS field is used by the switch to determine the type of service provided to the packet.

Observing DiffServ code points as a traffic grouping mechanism for defining QoS policies and overwriting the Diffserv code point fields are supported.

Figure 17 shows the encapsulation of an IP packet header.

Figure 17: IP packet header encapsulation

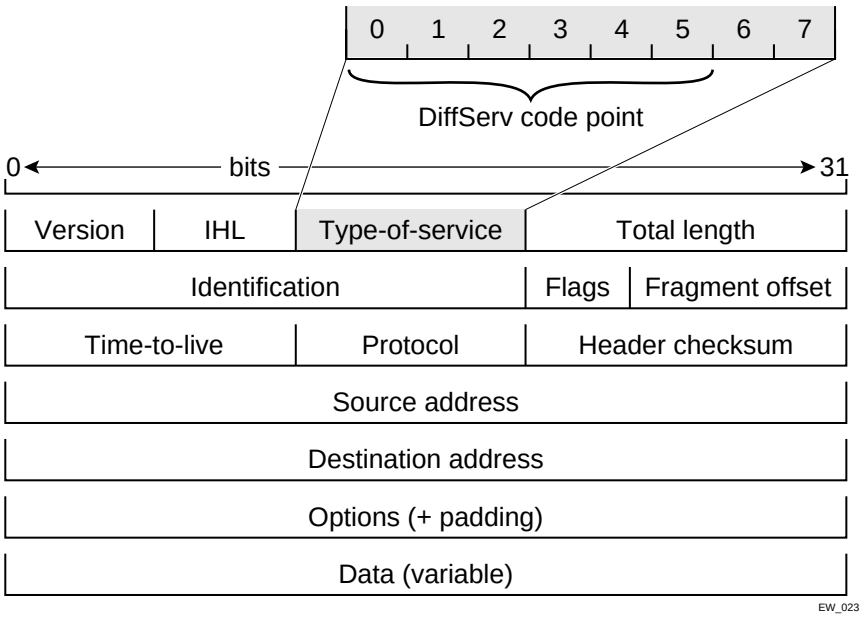

Observing DiffServ Information

When a packet arrives at the switch on an ingress port, the switch examines the first six of eight TOS bits, called the *code point*. The switch can assign the QoS profile used to subsequently transmit the packet based on the code point. The QoS profile controls a hardware queue used when transmitting the packet out of the switch, and determines the forwarding characteristics of a particular code point. Viewing DiffServ information can be enabled or disabled; by default it is disabled.

To enable DiffServ information, use the following command:

```
enable diffserv examination ports [<portlist> | all]
```

To disable DiffServ information, use the following command:

```
disable diffserv examination ports [<portlist> | all]
```

NOTE

After DiffServ information is enabled, the ACL router cannot apply for the same port.

Changing DiffServ Code point assignments in the Q0S Profile

Because the code point uses six bits, it has 64 possible values ($2^6 = 64$). Be default, the values are grouped and assigned to the default QoS profiles listed in Table 23.

Table 23: Default Code Point-to-QoS Profile Mapping

Code Point	QoS Profile
0-7	Qp1

Table 23: Default Code Point-to-QoS Profile Mapping (Continued)

Code Point	QoS Profile
8-15	Qp2
16-23	Qp3
24-31	Qp4
32-39	Qp5
40-47	Qp6
48-55	Qp7
56-63	Qp8

Once assigned, the rest of the switches in the network prioritize the packet using the characteristics specified by the QoS profile.

Replacing DiffServ Code Points

An access list can be used to change the DiffServ code point in the packet prior to the packet being transmitted by the switch. This is done with no impact on switch performance.

To replace the DiffServ code point, you will use an access list to set the new code point value. See “IP Access Lists (ACLs)” on page 138, for more information on using access lists. Use the `set code-point` parameter of the `create access-list` command to replace the value.

To display the DiffServ configuration, use the following command:

```
show ports {mgmt | <portlist>} info {detail}
```

NOTE

The show ports command displays only the default code point mapping.

Physical and Logical Groupings

Two traffic groupings exist in this category:

- Source port
- VLAN

Source port

A source port traffic grouping implies that any traffic sourced from this physical port uses the indicated QoS profile when the traffic is transmitted out to any other port. To configure a source port traffic grouping, use the following command:

```
configure ports <portlist> qosprofile <qosprofile>
```

In the following example, all traffic sourced from port 7 uses the QoS profile named *qp3* when being transmitted.

```
configure ports 7 qosprofile qp3
```

VLAN

A VLAN traffic grouping indicates that all intra-VLAN switched traffic and all routed traffic sourced from the named VLAN uses the indicated QoS profile. To configure a VLAN traffic grouping, use the following command:

```
configure vlan <vlan name> qosprofile <qosprofile>
```

For example, all devices on VLAN *servnet* require use of the QoS profile *qp4*. The command to configure this example is as follows:

```
configure vlan servnet qosprofile qp4
```

Verifying Physical and Logical Groupings

To verify settings on ports or VLANs, use the following command:

```
show qosprofile
```

The same information is also available for ports or VLANs using one of the following command:

```
show vlan
```

Verifying Configuration and Performance

Once you have created QoS policies that manage the traffic through the switch, you can use the QoS monitor to determine whether the application performance meets your expectations.

QoS Monitor

The QoS monitor is a utility that monitors the eight hardware queues (QP1-QP8) associated with any port(s). The QoS monitor keeps track of the number of frames and the frames per second that a specific ingress queue is responsible for transmitting on a physical port. Two options are available: a real-time display, and a separate option for retrieving information in the background and writing it to the log.

Real-Time Performance Monitoring

The real-time display scrolls through the given portlist to provide statistics. You can choose screens for packet count and packets per second. The specific port being monitored is indicated by an asterisk (*) appearing after the port number in the display.

The view real-time switch per-port performance, use the following command:

```
show ports {mgmt | <portlist>} qosmonitor
```

QoS monitor sampling is configured as follows:

- The port is monitored for 20 seconds before the switch moves on to the next port in the list.
- A port is sampled for five seconds before the packets per second (pps) value is displayed on the screen.

NOTE

The QoS monitor displays the statistics of incoming packets. The real-time display corresponds to the 802.1p values of the incoming packets. Any priority changes within the switch are not reflected in the display.

Displaying QoS Profile Information

The QoS monitor can also be used to verify the QoS configuration and monitor the use of the QoS policies that are in place. To display QoS information on the switch, use the following command:

```
show qosprofile {<qosprofile>} {port <portlist>}
```

Displayed information includes:

- QoS profile name
- Priority
- A list of all ports to which the QoS profile is applied
- A list of all VLANs to which the QoS profile is applied

Additionally, QoS information can be displayed from the traffic grouping perspective by using one or more of the following commands:

- `show fdb permanent`—Displays destination MAC entries and their QoS profiles.
- `show switch`—Displays hardware information.
- `show vlan`—Displays the QoS profile assignments to the VLAN.

Modifying a QoS Configuration

If you make a change to the parameters of a QoS profile after implementing your configuration, the timing of the configuration change depends on the traffic grouping involved. The following rules apply:

- For destination MAC-based grouping (other than permanent), clear the MAC FDB using the command `clear fdb`. This command should also be issued after a configuration is implemented, as the configuration must be in place before an entry is made in the MAC FDB. For permanent destination MAC-based grouping, re-apply the QoS profile to the static FDB entry, as documented. You can also save and reboot the switch.
- For physical and logical groupings of a source port or VLAN, re-apply the QoS profile to the source port or VLAN, as documented. You can also save and reboot the switch.

Traffic Rate-Limiting

The Summit 400 switch rate-limiting method is based on creating a rate limit, a specific type of access control list. Traffic that matches a rate limit is constrained to the limit set in the access control list. Rate limits are discussed in “Rate Limits” on page 139.

8 Status Monitoring and Statistics

This chapter describes the following topics:

Viewing statistics on a regular basis allows you to see how well your network is performing. If you keep simple daily records, you will see trends emerging and notice problems arising before they cause major network faults. In this way, statistics can help you get the best out of your network.

Port Statistics

ExtremeWare provides a facility for viewing port statistic information. The summary information lists values for the current counter against each port on each operational module in the system, and it is refreshed approximately every 2 seconds. Values are displayed to nine digits of accuracy.

To view port statistics, use the following command:

```
show ports {mgmt | <portlist>} stats
```

The following port statistic information is collected by the switch:

- **Link Status—**The current status of the link. Options are:
  - Ready (the port is ready to accept a link).
  - Active (the link is present at this port).
  - Chassis (the link is connected to a Summit Virtual Chassis).
- **Transmitted Packet Count (Tx Pkt Count)—**The number of packets that have been successfully transmitted by the port.
- **Transmitted Byte Count (Tx Byte Count)—**The total number of data bytes successfully transmitted by the port.

- **Received Packet Count (Rx Pkt Count)—**The total number of good packets that have been received by the port.
- **Received Byte Count (RX Byte Count)—**The total number of bytes that were received by the port, including bad or lost frames. This number includes bytes contained in the Frame Check Sequence (FCS), but excludes bytes in the preamble.
- **Received Broadcast (RX Bcast)—**The total number of frames received by the port that are addressed to a broadcast address.
- **Received Multicast (RX Mcast)—**The total number of frames received by the port that are addressed to a multicast address.

Port Errors

The switch keeps track of errors for each port.

To view port transmit errors, use the following command:

```
show ports {mgmt | <portlist>} txerrors
```

The following port transmit error information is collected by the system:

- **Port Number**
- **Link Status—**The current status of the link. Options are:
  - Ready (the port is ready to accept a link).
  - Active (the link is present at this port).
- **Transmit Collisions (TX Coll)—**The total number of collisions seen by the port, regardless of whether a device connected to the port participated in any of the collisions.
- **Transmit Late Collisions (TX Late Coll)—**The total number of collisions that have occurred after the port's transmit window has expired.
- **Transmit Deferred Frames (TX Deferred)—**The total number of frames that were transmitted by the port after the first transmission attempt was deferred by other network traffic.
- **Transmit Errored Frames (TX Error)—**The total number of frames that were not completely transmitted by the port because of network errors (such as late collisions or excessive collisions).
- **Transmit Parity Frames (TX Parity)**—The bit summation has a parity mismatch.

To view port receive errors, use the following command:

```
show ports {mgmt | <portlist>} rxerrors
```

The following port receive error information is collected by the switch:

- **Receive Bad CRC Frames (RX CRC)—**The total number of frames received by the port that were of the correct length, but contained a bad FCS value.
- **Receive Oversize Frames (RX Over)—**The total number of good frames received by the port greater than the supported maximum length of 1,522 bytes. For products that use the "*i*" chipset, ports with jumbo frames enabled do not increment this counter.
- **Receive Undersize Frames (RX Under)—**The total number of frames received by the port that were less than 64 bytes long.
- **Receive Fragmented Frames (RX Frag)—**The total number of frames received by the port were of incorrect length and contained a bad FCS value.

- **Receive Jabber Frames (RX Jab)—**The total number of frames received by the port that was of greater than the support maximum length and had a Cyclic Redundancy Check (CRC) error.
- **Receive Alignment Errors (RX Align)—**The total number of frames received by the port that occurs if a frame has a CRC error and does not contain an integral number of octets.
- **Receive Frames Lost (RX Lost)—**The total number of frames received by the port that were lost because of buffer overflow in the switch.

Port Monitoring Display Keys

Table 24 describes the keys used to control the displays that appear when you issue any of the `show port` commands.

Table 24: Port Monitoring Display Keys

Key(s)	Description
U	Displays the previous page of ports.
D	Displays the next page of ports.
[Esc] or [Return]	Exits from the screen.
0	Clears all counters.
[Space]	Cycles through the following screens: • Packets per second • Bytes per second • Percentage of bandwidth Available using the `show port utilization` command only.

Setting the System Recovery Level

You can configure the system to automatically reboot after a software task exception, using the following command:

```
configure sys-recovery-level [none | [all | critical] [ reboot | shutdown |
system-dump [maintenance-mode | reboot | shutdown]]]
```

Where the following is true:

- none—Configures the level to no recovery.
- `all`—Configures ExtremeWare to log an error into the syslog and automatically reboot the system after any task exception.
- `critical`—Configures ExtremeWare to log an error into the syslog and automatically reboot the system after a critical task exception.

The default setting is none.

Event Management System/Logging

Beginning in ExtremeWare 7.1.0, the system responsible for logging and debugging was updated and enhanced. We use the general term, event, for any type of occurrence on a switch which could generate a log message, or require an action. For example, a link going down, a user logging in, a command entered on the command line, or the software executing a debugging statement, are all events that might generate a log message. The new system for saving, displaying, and filtering events is called the Event Management System (EMS). With EMS, you have a lot more options about which events generate log messages, where the messages are sent, and how they are displayed. Using EMS you can:

- send event messages to a number of logging targets (for example, syslog host and NVRAM)
- filter events on a per-target basis
  - by component, subcomponent, or specific condition (for example, *IGMP.Snooping* messages, or the *IP.Forwarding.SlowPathDrop* condition)
  - by match expression (for example, any messages containing the string “user5”)
  - by matching parameters (for example, only messages with source IP addresses in the 10.1.2.0/24 subnet)
  - by severity level (for example, only messages of severity critical, error, or warning)
- change the format of event messages (for example, display the date as “12-May-2003” or “2003-05-12”)
- display log messages in real-time, and filter the messages that are displayed, both on the console and from telnet sessions
- display stored log messages from the memory buffer or NVRAM
- upload event logs stored in memory to a TFTP server
- display counts of event occurrences, even those not included in filter
- display debug information, using a consistent configuration method

Sending Event Messages to Log Targets

There are five types of targets that can receive log messages:

- console display
- current session (telnet or console display)
- memory buffer (can contain 200-20,000 messages)
- NVRAM (messages remain after reboot)
- syslog host

The first four types of targets exist by default, but before enabling any syslog host, the host’s information needs to be added to the switch using the `configure syslog` command. Extreme Networks EPICenter can be a syslog target.

By default, the memory buffer and NVRAM targets are already enabled and receive messages. To start sending messages to the targets, use the following command:

```
enable log target [console-display | memory-buffer | nvram | session | syslog [<host
name/ip> {:<udp-port>} [local0 ... local7]]]
```

Once enabled, the target receives the messages it is configured for. See the section “Target Configuration” for information on viewing the current configuration of a target. The memory buffer can only contain the configured number of messages, so the oldest message is lost when a new message arrives, and the buffer is full.

Use the following command to stop sending messages to the target:

```
disable log target [console-display | memory-buffer | nvram | session | syslog
[<host name/ip> {:<udp-port>} [local0 ... local7]]]
```

NOTE

Refer to your UNIX documentation for more information about the syslog host facility.

Filtering Events Sent to Targets

Not all event messages are sent to every enabled target. Each target receives only the messages that it is configured for.

Target Configuration

To specify the messages to send to a enabled target, you will set a message severity level, a filter name, and a match expression. These items determine which messages are sent to the target. You can also configure the format of the messages in the targets. Each target has a default configuration that mimics the expected behavior of prior ExtremeWare releases. For example, the console display target is configured to get messages of severity `info` and greater, the NVRAM target gets messages of severity `warning` and greater, and the memory buffer target gets messages of severity `debug-data` and greater. All the targets are associated by default with a filter named *DefaultFilter*, that passes all events at or above the default severity threshold, like the behavior of earlier releases (the earlier releases had no filters). All the targets are also associated with a default match expression that matches any messages (the expression that matches any message is displayed as `Match: (none)` from the command line). And finally, each target has a format associated with it.

To display the current log configuration of the targets, use the following command:

```
show log configuration target {console-display | memory-buffer | nvram | session |
syslog <host name/ip> {: <udp-port>}[local0 ... local7]}
```

To configure a target, there are specific commands for filters, formats, and severity that are discussed in the following sections.

Severity

Messages are issued with one of the severity level specified by the standard BSD syslog values (RFC 3164), `critical`, `error`, `warning`, `notice`, and `info`, plus three severity levels for extended debugging, `debug-summary`, `debug-verbose`, and `debug-data`. Note that RFC 3164 syslog values `emergency` and `alert` are not needed since `critical` is the most severe event in the system.

The three severity levels for extended debugging, `debug-summary`, `debug-verbose`, and `debug-data`, require that debug mode be enabled (which may cause a performance degradation). See the section “Displaying Debug Information” for more information about debugging.

Table 25: Severity Levels Assigned by the Switch[1]

Level	Description
Critical	A serious problem has been detected which is compromising the operation of the system and that the system can not function as expected unless the situation is remedied. The switch may need to be reset.
Error	A problem has been detected which is interfering with the normal operation of the system and that the system is not functioning as expected.
Warning	An abnormal condition, not interfering with the normal operation of the system, has been detected which may indicate that the system or the network in general may not be functioning as expected.
Notice	A normal but significant condition has been detected, which signals that the system is functioning as expected.
Info (Informational)	A normal but potentially interesting condition has been detected, which signals that the system is functioning as expected and simply provides potentially detailed information or confirmation.
Debug-Summary	A condition has been detected that may interest a developer determining the reason underlying some system behavior.
Debug-Verbose	A condition has been detected that may interest a developer analyzing some system behavior at a more verbose level than provided by the debug summary information.
Debug-Data	A condition has been detected that may interest a developer inspecting the data underlying some system behavior.

1. In ExtremeWare version 7.1.0, the levels alert and emergency were deprecated. The equivalent level is critical.

To configure the severity level of the messages sent to a target, there is more than one command that you can use. The most direct way to set the severity level of all the sent messages is to use the following command:

```
configure log target [console-display | memory-buffer | nvram | session |
syslog [<host name/ip> {: <udp-port>} [local0 ... local7]]]
filter <filter name> {severity <severity> {only}}
```

When you specify a severity level, messages of that severity and greater will be sent to the target. If you want only messages of the specified severity to be sent to the target, use the keyword `only`. For example, specifying `severity warning` will send warning, error, and critical messages, but specifying `severity warning only` will just send warning messages.

Another command that can be used to configure severity levels is the command used to associate a filter with a target:

```
configure log target [console-display | memory-buffer | nvram | session |
syslog [<host name/ip> {: <udp-port>} [local0 ... local7]]]
filter <filter name> {severity <severity> {only}}
```

When you specify a severity level as you associate a filter with a target, you further restrict the messages reaching the target. The filter may only allow certain categories of messages to pass. Only the messages that pass the filter, and then pass the specified severity level will reach the target.

Finally, you can specify the severity levels of messages that reach the target by associating a filter with a target. The filter can specify exactly which message it will pass. Constructing a filter is discussed in the section “Filtering By Components and Conditions”.

Components and Conditions

Beginning with the introduction of EMS in release 7.1.0, the event conditions detected by ExtremeWare were organized into components and subcomponents. This is somewhat similar to the fault log subsystems used in previous versions. Not all conditions have been placed in the component/subcomponent structure of EMS, but all the conditions will be moved over time into this structure. To get a listing of the components and subcomponents in your release of ExtremeWare, use the following command:

```
show log components {<event component> | all}
```

For example, to get a listing of the subcomponents that make up the STP component, use the following command:

```
show log components stp
```

The output produced by the command is similar to the following:

```
                                                                  Severity
Component           Title                                         Threshold
------------------- --------------------------------------------- ----------
STP                 Spanning-Tree Protocol (STP)                  Error
        InBPDU      STP In BPDU subcomponent                      Warning
        OutBPDU     STP Out BPDU subcomponent                     Warning
        System      STP System subcomponent                       Error
```

In the display above is listed the component, the subcomponents that make up that component, and the default severity threshold assigned to that component. A period (.) is used to separate component, subcomponent, and condition names in EMS. For example, you can refer to the *InBPDU* subcomponent of the *STP* component as *STP.InBPDU*. On the CLI, you can abbreviate or TAB complete any of these.

A component or subcomponent will often have several conditions associated with it. To see the conditions associated with a component, use the following command:

```
show log events {<event condition> | [all | <event component>] {severity <severity>
{only}}} {detail}
```

For example, to see the conditions associated with the *STP.InBPDU* subcomponent, use the following command:

```
show log events stp.inbpdu
```

The output produced by the command is similar to the following:

```
Comp    SubComp     Condition              Severity      Parameters
------- ----------- ---------------------- ------------ ----------
STP     InBPDU
                    Drop                   Error         3
                    Dump                   Debug-Data    3
                    Ign                    Debug-Summary 2
                    Trace                  Info          2
```

In the display above is listed the four conditions contained in the *STP.InBPDU* component, the severity of the condition, and the number of parameters in the event message. In this example, the severities of the events in the *STP.InBPDU* subcomponent range from error to debug-summary.

When you use the `detail` keyword you will see the message text associated with the conditions. For example, if you want to see the message text and the parameters for the event condition *STP.InBPDU.Trace*, use the following command:

```
show log events stp.inbpdu.trace detail
```

The output produced by the command is similar to the following:

```
Comp    SubComp     Condition               Severity      Parameters
------- ----------- ----------------------- ------------ ----------
STP     InBPDU      Trace                   Info          2 Total
                                                          0 - ports
                                                          1 - string
                    "Port=%0%: %1%"
```

The `Comp` heading shows the component name, the `SubComp` heading shows the subcomponent (if any), the `Condition` heading shows the event condition, the `Severity` heading shows the severity assigned to this condition, the `Parameters` heading shows the parameters for the condition, and the text string shows the message that the condition will generate. The parameters in the text string (for example,`%0%` and`%1%` above) will be replaced by the values of these parameters when the condition is encountered, and output as the event message.

Filtering By Components and Conditions. You may want to send the messages that come from a specific component that makes up ExtremeWare, or send the message generated by a specific condition. For example, you might want to send only the messages that come from the STP component, or send the message that occurs when the *IP.Forwarding.SlowPathDrop* condition occurs. Or you may want to exclude messages from a particular component or event. To do this, you will construct a filter that passes only the items of interest, and associate that filter with a target.

The first step is to create the filter using the `create log filter` command. You can create a filter from scratch, or copy another filter to use as a starting point. It may be easiest to copy an existing filter and modify it. Use the following command to create a filter:

```
create log filter <name> {copy <filter name>}
```

If you create a filter from scratch, it will initially block all events until you add events (either the events from a component or a specific event condition) to pass. You might create a filter from scratch if you wanted to pass a small set of events, and block most. If you want to exclude a small set of events, there is a default filter that passes events at or above the default severity threshold (unless the filter has been modified), named *DefaultFilter*, that you can copy to use as a starting point for your filter.

Once you have created your filter, you can then configure filter items that include or exclude events from the filter. Included events are passed, excluded events are blocked. Use the following command to configure your filter:

```
configure log filter <filter name> [add | delete] {exclude} events [<event condition>
| [all | <event component>] {severity <severity> {only}}]
```

For example, if you create the filter *myFilter* from scratch, then issue the following command:

```
configure log filter myFilter add events stp
```

all STP events will pass *myFilter* of at least the default threshold severity (for the STP component, the default severity threshold is `error`). You can further modify this filter by specifying additional conditions. For example, assume that *myFilter* is configured as before, and assume that you want to exclude any events from the STP subcomponent, *STP.OutBPDU*. Use the following command to add that condition:

```
configure log filter myFilter add exclude events stp.outbpdu
```

You can continue to modify this filter by adding more filter items. The filters process events by comparing the event with the most recently configured filter item first. If the event matches this filter

item, the incident is either included or excluded, depending on whether the exclude keyword was used. Subsequent filter items on the list are compared if necessary. If the list of filter items has been exhausted with no match, the event is excluded, and is blocked by the filter.

To examine the configuration of a filter, use the following command:

```
show log configuration filter {<filter name>}
```

The output produced by the command (for the earlier filter) is similar to the following:

```
Log Filter Name : myFilter
I/                                              Severity
E  Comp    SubComp     Condition                CEWNISVD
-  ------- ----------- ------------------------ --------
E  STP     OutBPDU     *                        CEWNI+++
I  STP     *           *                        ********

Include/Exclude: (I) Include,  (E) Exclude
Severity Values: (C) Critical,  (E) Error,  (W) Warning,  (N) Notice,  (I) Info
                 (*) Pre-assigned severities in effect for each subcomponent
Debug Severity : (S) Debug-Summary,  (V) Debug-Verbose,  (D) Debug-Data
                 (+) Debug Severity requested, but log debug-mode not enabled
If Match parameters present:
Parameter Flags: (S) Source,  (D) Destination  (as applicable)
                 (I) Ingress,  (E) Egress,
Parameter Types: Port - Physical Port list,  Slot - Physical Slot #
                 MAC  - MAC address,  IP - IP Address/netmask,  Mask - Netmask
                 VID  - Virtual LAN ID (tag),  VLAN - Virtual LAN name
                   Nbr  - Neighbor, Rtr  - Routerid, EAPS - EAPS Domain
Strict Match   : (Y) every match parameter entered must be present in the event
                 (N) match parameters need not be present in the event
```

The show log configuration filter command shows each filter item, in the order that it will be applied and whether it will be included or excluded. The above output shows the two filter items, one excluding events from the *STP.OutBPDU* component, the next including the remaining events from the *STP* component. The severity value is shown as "*", indicating that the component's default severity threshold controls which messages are passed. The Parameter(s) heading is empty for this filter, since no match was configured for this filter. Matches are discussed in the section, "Matching Expressions".

Each time a filter item is added to or deleted from a given filter, the events specified are compared against the current configuration of the filter to try to logically simplify the configuration. Existing items will be replaced by logically simpler items if the new item enables rewriting the filter. If the new item is already included or excluded from the currently configured filter, the new item is not added to the filter.

Matching Expressions

You can specify that messages that reach the target match a specified match expression. The message text is compared with the match expression to determine whether to pass the message on. To require that messages match a match expression, is to use the following command:

```
configure log target [console-display | memory-buffer | nvram | session |
syslog [<host name/ip> {: <udp-port>} [local0 ... local7]]] match [any
|<match-expression>]
```

The messages reaching the target will match the match-expression, a simple regular expression. The formatted text string that makes up the message is compared with the match expression, and is passed to the target if it matches. This command does not affect the filter in place for the target, so the match

expression is only compared with the messages that have already passed the target's filter. For more information on controlling the format of the messages, see the section, "Formatting Event Messages".

Simple Regular Expressions. A simple regular expression is a string of single characters including the dot character (.), which are optionally combined with quantifiers and constraints. A dot matches any single character while other characters match only themselves (case is significant). Quantifiers include the star character (*) that matches zero or more occurrences of the immediately preceding token. Constraints include the caret character (^) that matches at the beginning of a message, and the currency character ($) that matches at the end of a message. Bracket expressions are not supported. There are a number of sources available on the Internet and in various language references describing the operation of regular expressions. Table 26 shows some examples of regular expressions.

Table 26: Simple Regular Expressions

Regular Expression	Matches	Does not match
port	port 2:3 import cars portable structure	poor por pot
.ar	baar bazaar rebar	bar
port.*vlan	port 2:3 in vlan test add ports to vlan port/vlan	
myvlan$	delete myvlan error in myvlan	myvlan port 2:3 ports 2:4,3:4 myvlan link down

Matching Parameters

Rather than using a text match, ExtremeWare's EMS allows you to filter more efficiently based on the message parameter values. In addition to event components and conditions and severity levels, each filter item can also use parameter values to further limit which messages are passed or blocked. The process of creating, configuring, and using filters has already been described in the section, "Filtering By Components and Conditions", so this section will discuss matching parameters with a filter item. To configure a parameter match filter item, use the following command:

```
configure log target [console-display | memory-buffer | nvram | session |
syslog [<host name/ip> {: <udp-port>} [local0 ... local7]]]
filter <filter name> {severity <severity> {only}}
```

Each event in ExtremeWare is defined with a message format and zero or more parameter types. The `show log events detail` command can be used to display event definitions (the event text and parameter types). Only those parameter types that are applicable given the events and severity specified are exposed on the CLI. The <value> depends on the parameter type specified. As an example, an event may contain a physical port number, a source MAC address, and a destination MAC address. To allow only those Bridging incidents, of severity `notice` and above, with a specific source MAC address, use the following command:

```
configure log filter myFilter add events bridge severity notice match source
mac-address 00:01:30:23:C1:00
```

The string type is used to match a specific string value of an event parameter, such as a user name. A string can be specified as a simple regular expression.

Use the and keyword to specify multiple parameter type/value pairs that must match those in the incident. For example, to allow only those events with specific source and destination MAC addresses, use the following command:

```
configure log filter myFilter add events bridge severity notice match source
mac-address 00:01:30:23:C1:00 and destination mac-address 01:80:C2:00:00:02
```

Match Versus Strict-Match. The `match and strict-match` keywords control the filter behavior for incidents whose event definition does not contain all the parameters specified in a `configure log filter events match` command. This is best explained with an example. Suppose an event in the *XYZ* component, named *XYZ.event5*, contains a physical port number, a source MAC address, but no destination MAC address. If you configure a filter to match a source MAC address and a destination MAC address, *XYZ.event5* will match the filter when the source MAC address matches regardless of the destination MAC address, since the event contains no destination MAC address. If you specify the `strict-match` keyword, then the filter will never match event *XYZ.event5*, since this event does not contain the destination MAC address.

In other words, if the `match` keyword is specified, an incident will pass a filter so long as all parameter values in the incident match those in the match criteria, but all parameter types in the match criteria need not be present in the event definition.

Formatting Event Messages

Event messages are made up of a number of items. The individual items can be formatted, however, EMS does not allow you to vary the order of the items. To format the messages for a particular target, use the following command:

```
configure log target [console-display | memory-buffer | nvram | session |
syslog [<host name/ip> {:<udp-port>} [local0 ... local7]]]
format [timestamp [seconds | hundredths | none]
| date [dd-mm-yyyy | dd-Mmm-yyyy | mm-dd-yyyy | Mmm-dd | yyyy-mm-dd | none]
| severity [on | off]
| event-name [component | condition | none | subcomponent]
| host-name [on | off]
| priority [on | off]
| tag-id [on | off]
| tag-name [on | off]
| sequence-number [on | off]
| process-name [on | off]
| process-id [on | off]
| source-function [on | off]
| source-line [on | off]]
```

Using the default format for the session target, an example log message might appear as:

```
05/29/2003 12:15:25.00 <Warn:SNTP.RslvSrvrFail> The SNTP server parameter value
(TheWrongServer.example.com) can not be resolved.
```

If you set the current session format using the following command:

```
configure log target session format date mm-dd-yyy timestamp seconds event-name
component
```

The same example would appear as:

```
05/29/2003 12:16:36 <Warn:SNTP> The SNTP server parameter value
(TheWrongServer.example.com) can not be resolved.
```

In order to provide some detailed information to technical support, you set the current session format using the following command:

```
configure log target session format date mmm-dd timestamp hundredths event-name
condition source-line on process-name on
```

The same example would appear as:

```
May 29 12:17:20.11 SNTP: <Warn:SNTP.RslvSrvrFail> tSntpc: (sntpcLib.c:606) The SNTP
server parameter value (TheWrongServer.example.com) can not be resolved.
```

Displaying Real-Time Log Messages

You can configure the system to maintain a running real-time display of log messages on the console display or on a (telnet) session. To turn on the log display on the console, use the `console-display` option in the following command:

```
enable log target [console-display | memory-buffer | nvram | session | syslog [<host
name/ip> {:<udp-port>} [local0 ... local7]]]
```

This setting may be saved to the FLASH configuration and will be restored on boot up (to the console-display session).

To turn on log display for the current session:

```
enable log target session
```

This setting only affects the current session, and is lost when you log off the session.

The messages that are displayed depend on the configuration and format of the target. See the section, “Filtering Events Sent to Targets”, for information on message filtering, and the section, “Formatting Event Messages”, for information on message formatting.

Displaying Events Logs

The log stored in the memory buffer and the NVRAM can be displayed on the current session (either the console display or telnet). Use the following command to display the log:

```
show log {messages [memory-buffer | nvram]} {severity <severity> {only}}
{starting [date <date> time <time> | date <date> | time <time>]} {ending [date
<date> time <time> | date <date> | time <time>]} {match <match-expression>}
{format <format>} {chronological}
```

There are many options you can use to select the log entries of interest. You can select to display only those messages that conform to the specified:

- severity
- starting and ending date and time
- match expression

The displayed messages can be formatted differently from the format configured for the targets, and you can choose to display the messages in order of newest to oldest, or in chronological order (oldest to newest).

Uploading Events Logs

The log stored in the memory buffer and the NVRAM can be uploaded to a TFTP server. Use the following command to upload the log:

```
upload log <host name/ip> <filename> {messages [memory-buffer | nvram]}
{severity <severity> {only}} {starting [date <date> time <time> | date <date>
| time <time>]} {ending [date <date> time <time> | date <date> | time <time>]}
{match <match-expression>} {format <format>} {chronological}
```

You must specify the TFTP host and the filename to use in uploading the log. There are many options you can use to select the log entries of interest. You can select to upload only those messages that conform to the specified:

- severity
- starting and ending date and time
- match expression

The uploaded messages can be formatted differently from the format configured for the targets, and you can choose to upload the messages in order of newest to oldest, or in chronological order (oldest to newest).

Displaying Counts of Event Occurrences

EMS adds the ability to count the number of occurrences of events. Even when an event is filtered from all log targets, the event is counted. (The exception to this is events of any of the debug severities, which are only counted when the log debug mode is enabled.) To display the event counters, use the following command:

```
show log counters {<event condition> | [all | <event component>] {severity <severity>
{only}}}
```

Two counters are displayed. One counter displays the number of times an event has occurred, and the other displays the number of times that notification for the event was made to the system for further processing. Both counters reflect totals accumulated since reboot or since the counters were cleared using the `clear log counters` or `clear counters` command.

This command also displays an included count (the column titled `In` in the output). The reference count is the number of enabled targets receiving notifications of this event without regard to matching parameters.

The keywords `included`, `notified`, and `occurred` only display events with non-zero counter values for the corresponding counter.

Output of the command:

```
show log counters stp.inbpdu severity debug-summary
```

will be similar to the following:

```
Comp     SubComp     Condition               Severity      Occurred  In Notified
-------- ----------- ----------------------- ------------- --------  -- --------
STP      InBPDU
                     Drop                    Error                0   1        0
                     Ign                     Debug-Summary        0+  0        0
```

```
                    Trace                    Info                 0   0        0
Occurred  : # of times this event has occurred since last clear or reboot
Flags     : (+) Debug events are not counted while log debug-mode is disabled
In(cluded): # of enabled targets whose filter includes this event
Notified  : # of times this event has occurred when 'Included' was non-zero
```

Output of the command:

```
show log counters stp.inbpdu.drop
```

will be similar to the following:

```
Comp     SubComp     Condition               Severity      Occurred  In Notified
-------- ----------- ----------------------- ------------ --------  -- --------
STP      InBPDU
                     Drop                    Error                0   1        0
```

Displaying Debug Information

By default, a switch will not generate events of severity `Debug-Summary`, `Debug-Verbose`, and `Debug-Data` unless the switch is in debug mode. Debug mode causes a performance penalty, so it should only be enabled for specific cases where it is needed. To place the switch in debug mode, use the following command:

```
enable log debug-mode
```

Once debug mode is enabled, any filters configured for your targets will still affect which messages are passed on or blocked.

NOTE

Previous versions of ExtremeWare used the debug-trace *command to enable debugging. Not all systems in ExtremeWare were converted to use EMS in the initial release. As a result, some debug information still requires you to use the corresponding* debug-trace *command. The* show log component *command displays the systems in your image that are part of EMS. Any systems in EMS will not have* debug-trace *commands, and vice-versa*

Compatibility with previous ExtremeWare commands

Since EMS provides much more functionality, there are a number of new commands introduced to support it. However, if you do not require the enhanced capabilities provided by EMS, you can continue to use many of the logging commands that existed in earlier versions of ExtremeWare. For consistency, the earlier commands are still supported. Listed below are earlier commands with their new command equivalents.

Enable / disable log display

The following commands related to the serial port console:

```
enable log display
disable log display
```

are equivalent to using the console-display option in the following commands:

```
enable log target [console-display | memory-buffer | nvram | session | syslog [<host
name/ip> {:<udp-port>} [local0 ... local7]]]

disable log target [console-display | memory-buffer | nvram | session | syslog
[<host name/ip> {:<udp-port>} [local0 ... local7]]]
```

Note that the existing command `enable log display` applies only to the serial port console. Since the ability to display log messages on other sessions was added, the target name `session` was chosen. For clarity, the target name `console-display` was chosen to refer to the serial port console, previously referred to as simply `display`.

Configure log display

The following command related to the serial port console:

```
configure log display {<severity>}
```

is equivalent to:

```
configure log target console-display severity <severity>
```

Remote syslog commands

The following command related to remote syslog hosts:

```
configure syslog {add} <host name/ip> {: <udp-port>} [local0 ... local7] {<severity>}
```

is equivalent to the following two commands:

```
configure syslog add <hostname/IP> {: <udp-port>} [local0 ... local7]
configure log target syslog <hostname/IP> {: <udp-port>} [local0 ... local7] severity
<severity>
```

NOTE

Refer to your UNIX documentation for more information about the syslog host facility.

Logging Configuration Changes

ExtremeWare allows you to record all configuration changes and their sources that are made using the CLI by way of telnet or the local console. The changes cause events that are logged to the target logs. Each log entry includes the user account name that performed the change and the source IP address of the client (if telnet was used). Configuration logging applies only to commands that result in a configuration change. To enable configuration logging, use the following command:

```
enable cli-config-logging
```

To disable configuration logging, use the following command:

```
disable cli-config-logging
```

CLI configuration logging is enabled by default.

RMON

Using the Remote Monitoring (RMON) capabilities of the switch allows network administrators to improve system efficiency and reduce the load on the network.

The following sections explain more about the RMON concept and the RMON features supported by the switch.

NOTE

You can only use the RMON features of the system if you have an RMON management application, and have enabled RMON on the switch.

About RMON

RMON is the common abbreviation for the Remote Monitoring Management Information Base (MIB) system defined by the Internet Engineering Task Force (IETF) documents RFC 1271 and RFC 1757, which allows you to monitor LANs remotely.

A typical RMON setup consists of the following two components:

- **RMON probe**—An intelligent, remotely controlled device or software agent that continually collects statistics about a LAN segment or VLAN. The probe transfers the information to a management workstation on request, or when a predefined threshold is crossed.
- **Management workstation**—Communicates with the RMON probe and collects the statistics from it. The workstation does not have to be on the same network as the probe, and can manage the probe by in-band or out-of-band connections.

RMON Features of the Switch

The IETF defines nine groups of Ethernet RMON statistics. The switch supports the following four of these groups:

- Statistics
- History
- Alarms
- Events

This section describes these groups and discusses how they can be used.

Statistics

The RMON Ethernet Statistics group provides traffic and error statistics showing packets, bytes, broadcasts, multicasts, and errors on a LAN segment or VLAN.

Information from the Statistics group is used to detect changes in traffic and error patterns in critical areas of the network.

History

The History group provides historical views of network performance by taking periodic samples of the counters supplied by the Statistics group. The group features user-defined sample intervals and bucket counters for complete customization of trend analysis.

The group is useful for analysis of traffic patterns and trends on a LAN segment or VLAN, and to establish baseline information indicating normal operating parameters.

Alarms

The Alarms group provides a versatile, general mechanism for setting threshold and sampling intervals to generate events on any RMON variable. Both rising and falling thresholds are supported, and thresholds can be on the absolute value of a variable or its delta value. In addition, alarm thresholds can be autocalibrated or set manually.

Alarms inform you of a network performance problem and can trigger automated action responses through the Events group.

Events

The Events group creates entries in an event log and/or sends SNMP traps to the management workstation. An event is triggered by an RMON alarm. The action taken can be configured to ignore it, to log the event, to send an SNMP trap to the receivers listed in the trap receiver table, or to both log and send a trap. The RMON traps are defined in RFC 1757 for rising and falling thresholds.

Effective use of the Events group saves you time. Rather than having to watch real-time graphs for important occurrences, you can depend on the Event group for notification. Through the SNMP traps, events can trigger other actions, which provides a mechanism for an automated response to certain occurrences.

Configuring RMON

RMON requires one probe per LAN segment, and standalone RMON probes traditionally have been expensive. Therefore, Extreme's approach has been to build an inexpensive RMON probe into the agent of each system. This allows RMON to be widely deployed around the network without costing more than traditional network management. The switch accurately maintains RMON statistics at the maximum line rate of all of its ports.

For example, statistics can be related to individual ports. Also, because a probe must be able to see all traffic, a stand-alone probe must be attached to a nonsecure port. Implementing RMON in the switch means that all ports can have security features enabled.

To enable or disable the collection of RMON statistics on the switch, use one of the following commands:

```
enable rmon
disable rmon
```

By default, RMON is disabled. However, even in the disabled state, the switch responds to RMON queries and sets for alarms and events. By enabling RMON, the switch begins the processes necessary for collecting switch statistics.

Event Actions

The actions that you can define for each alarm are shown in Table 27.

Table 27: Event Actions

Action	High Threshold
No action	
Notify only	Send trap to all trap receivers.
Notify and log	Send trap; place entry in RMON log.

To be notified of events using SNMP traps, you must configure one or more trap receivers, as described in Chapter 2.

9 Security

This chapter describes the following topics:

Security Overview

Extreme Networks products incorporate a number of features designed to enhance the security of your network. No one feature can insure security, but by using a number of features in concert, you can substantially improve the security of your network. The features described in this chapter are part of an overall approach to network security

Network Access Security

Network access security features control devices accessing your network. In this category are the following features:

- MAC-Based VLANs
- IP Access Lists (ACLs)
- Network Login

MAC-Based VLANs

MAC-Based VLANs allow physical ports to be mapped to a VLAN based on the source MAC address learned in the FDB. This feature allows you to designate a set of ports that have their VLAN membership dynamically determined by the MAC address of the end station that plugs into the physical port. You can configure the source MAC address-to-VLAN mapping either offline or dynamically on the switch. For example, you could use this application for a roaming user who wants to connect to a network from a conference room. In each room, the user plugs into one of the designated ports on the switch and is mapped to the appropriate VLAN. Connectivity is maintained to the network with all of the benefits of the configured VLAN in terms of QoS, routing, and protocol support.

Detailed information about configuring and using MAC-based VLANs can be found in Chapter 5.

IP Access Lists (ACLs)

Each access control list (ACL) consists of an access mask that selects which fields of each incoming packet to examine, and a list of values to compare with the values found in the packet. Access masks can be shared multiple access control lists, using different lists of values to examine packets. The following sections describe how to use access control lists.

Access Masks

There are sixteen access masks available in the Summit 400-48, depending on which features are enabled on the switch. Each access mask is created with a unique name and defines a list of fields that will be examined by any access control list that uses that mask (and by any rate limit that uses the mask).

To create an access mask, use the following command:

```
create access-mask <access-mask name> {dest-mac} {source-mac} {vlan} {tos
|code-point} {ethertype} {ipprotocol} {dest-ip/<mask length>} {source-L4port |
{icmp-type} {icmp-code}} {permit-established} {egresport} {ports} {precedence
<number>}
```

You can also display or delete an access mask. To display information about an access mask, use the following command:

```
show access-mask {<name>}
```

To delete an access mask, use the following command:

```
delete access-mask <name>
```

Access Lists

Access control lists are used to perform packet filtering and forwarding decisions on incoming traffic. Each packet arriving on an ingress port is compared to the access list in sequential order and is either forwarded to a specified QoS profile or dropped. These forwarded packets can also be modified by changing the 802.1p value and/or the DiffServ code point. Using access lists has no impact on switch performance.

The Summit 400-48 supports up to 16 access lists. Each entry that makes up an access list contains a unique name and specifies a previously created access mask. The access list also includes a list of values

to compare with the incoming packets, and an action to take for packets that match. When you create an access list, you must specify a value for each of the fields that make up the access mask used by the list.

To create an access list, use the following command:

```
create access-list <name> access-mask <access-mask name> {dest-mac <dest_mac}
{source-mac <src_mac>} {vlan <name>} {ethertype [IP | ARP | <hex_value>]} {tos
<ip_precedence> | code-point <code_point>} {ipprotocol [tcp | udp | icmp | igmp |
<prococol_num>]} {dest-ip <dest_IP>/<mask length>} {dest-L4port <dest_port>}
{source-ip <src_IP>/<mask length>} {source-L4port <src_port> [permit {qosprofile
<qosprofile>} {set code-point <code_point>} {set dot1p <dot1p_value} |
permit-established | deny]
```

NOTE

The parameters of the create access list command must match identically to the parameters of the create access-mask. The order of the parameters is also important. If the parameter are out-of-order, many of the options become unavailable to the user.

For packets that match a particular access list, you can specify the following actions:

- **Deny**—Matching packets are not forwarded.
- **Permit-established**—Drop the packet if it would initiate a new TCP session (see, “The permit-established Keyword” on page 141).
- **Permit**—Forward the packet. You can send the packet to a particular QoS profile, and modify the packet’s 802.1p value and / or DiffServ code point.

If a packet matches more than one access list, the switch uses the following rules to govern the actions of the packet:

- If the actions specified by the matching ACLs do not conflict, all of the actions are carried out.
- If the actions conflict, the associated access mask precedence determines the course of action. The access list with the highest precedence access-mask prevails.

To display information about one or more access lists, use the following command:

```
show access-list {<name> | port <portlist>}
```

To delete an access list, use the following command:

```
delete access-list <name>
```

Rate Limits

Rate limits are almost identical to access control lists. Incoming packets that match a rate limit access control list are allowed as long as they do not exceed a pre-defined rate. Excess packets are either dropped, or modified by resetting their DiffServ code point.

Each entry that makes up a rate limit contains a unique name and specifies a previously created access mask. Like an access list, a rate limit includes a list of values to compare with the incoming packets and an action to take for packets that match. Additionally, a rate limit specifies an action to take when

matching packets arrive at a rate above the limit you set. When you create a rate limit, you must specify a value for each of the fields that make up the access mask used by the list.

To create a rate limit rule, use the following command:

```
create rate-limit <rule_name> access-mask <access-mask name> {dest-mac <dest_mac>}
{source-mac <scr_mac>} {vlan <name>} {ethertype [IP | ARP | <hex_value>]} {tos
<ip_precedence> | code-point <code_point>} {ipprotocol [tcp | udp | icmp | igmp |
<prococol_num>]} {dest-ip <dest_IP>/<mask length>} {dest-L4port <dest_port>}
{source-ip <src_IP>/<mask length>} {source-L4port <src_port> [permit {qosprofile
<qosprofile>} {set code-point <code_point>} {set dot1p <dot1p_value} limit
<rate_in_Mbps> {exceed-action [drop | set code-point <code_point>]}
```

NOTE

Unlike an access list, a rate limit can only be applied to a single port. Each port will have its own rate limit defined separately.

On a 100 Mbps port (100BASE-TX), you can configure the rate limit value in the range from 1 Mbps to 100 Mbps in 1 Mbps increments, which is to say, the rate limit value can be set at 1, 2, 3, 4 … 100 Mbps.

On a 1000 Mbps port (Gigabit Ethernet uplink port), you can configure the rate limit value in the range from 8 Mbps to 1000 Mbps in increments of 8 Mbps, which is to say the rate limit value can be set at 8, 16, 24, 32 … 1000 Mbps.

NOTE

The rate limit specified in the command line does not precisely match the actual rate limit imposed by the hardware, due to hardware constraints. See the release notes for the exact values of the actual rate limits, if required for your implementation.

For packets that match a particular list, and arrive at a rate below the limit, you can specify the following action:

- **Permit**—Forward the packet. You can send the packet to a particular QoS profile, and modify the packet's 802.1p value and/or DiffServ code point.

For packets that match a particular list and arrive at a rate that exceeds the limit, you can specify the following actions:

- **Drop**—Drop the packets. Excess packets are not forwarded.
- **Permit with rewrite**—Forward the packet, but modify the packet's DiffServ code point.

How Access Control Lists Work

When a packet arrives on an ingress port, the fields of the packet corresponding to an access mask are compared with the values specified by the associated access lists to determine a match.

It is possible that a packet will match more than one access control list. If the resulting actions of all the matches do not conflict, they will all be carried out. If there is a conflict, the actions of the access list using the higher precedence access mask are applied. When a match is found, the packet is processed. If the access list is of type deny, the packet is dropped. If the list is of type permit, the packet is

forwarded. A permit access list can also apply a QoS profile to the packet and modify the packet's 802.1p value and the DiffServ code point.

Access Mask Precedence Numbers

The access mask precedence number determines the order in which each rule is examined by the switch and is optional. Access control list entries are evaluated from highest precedence to lowest precedence. Precedence numbers range from 1 to 25,600, with the *number 1 having the highest precedence*, but an access mask *without* a precedence specified has a higher precedence than any access mask *with* a precedence specified. The first access mask defined without a specified precedence has the highest precedence. Subsequent masks without a specified precedence have a lower precedence, and so on.

Specifying a Default Rule

You can specify a default access control list to define the default access to the switch. You should use an access mask with a low precedence for the default rule access control list. If no other access control list entry is satisfied, the default rule is used to determine whether the packet is forwarded or dropped. If no default rule is specified, the default behavior is to forward the packet.

NOTE

If your default rule denies traffic, you should not apply this rule to the Summit 400-48 port used as a management port.

Once the default behavior of the access control list is established, you can create additional entries using precedence numbers.

The `permit-established` Keyword

The `permit-established` keyword is used to directionally control attempts to open a TCP session. Session initiation can be explicitly blocked using this keyword.

The permit-established keyword denies the access control list. Having a permit-established access control list blocks all traffic that matches the TCP source/destination, and has the SYN=1 and ACK=0 flags set.

Adding Access Mask, Access List, and Rate Limit Entries

Entries can be added to the access masks, access lists, and rate limits. To add an entry, you must supply a unique name using the `create` command, and supply a number of optional parameters. For access lists and rate limits, you must specify an access mask to use. To modify an existing entry, you must delete the entry and retype it, or create a new entry with a new unique name.

To add an access mask entry, use the following command:

```
create access-mask <name> ...
```

To add an access list entry, use the following command:

```
create access-list <name> ...
```

To add a rate limit entry, use the following command:

```
create rate-limit <name> ...
```

Maximum Entries

If you try to create an access mask when no more are available, the system will issue a warning message. Three access masks are constantly used by the system, leaving a maximum of 13 user-definable access masks. However, enabling some features causes the system to use additional access masks, reducing the number available.

For each of the following features that you enable, the system will use one access mask. When the feature is disabled, the mask will again be available. The features are:

- RIP
- IGMP or OSPF (both would share a single mask)
- DiffServ examination
- QoS monitor

The maximum number of access list allowed by the hardware is 254 for each block of eight 10/100 Mbps Ethernet ports and 126 for each Gbps Ethernet port, for a total of 1014 rules (254*3+126*2). Most user entered access list commands will require multiple rules on the hardware. For example, a global rule (an access control list using an access mask without “ports” defined), will require 5 rules, one for each of the 5 blocks of ports on the hardware.

The maximum number of rate-limiting rules allowed is 315 (63*5). This number is part of the total access control list rules (1014).

Deleting Access Mask, Access List, and Rate Limit Entries

Entries can be deleted from access masks, access lists, and rate limits. An access mask entry cannot be deleted until all the access lists and rate limits that reference it are also deleted.

To delete an access mask entry, use the following command:

```
delete access-mask <name>
```

To delete an access list entry, use the following command:

```
delete access-list <name>
```

To delete a rate limit entry, use the following command:

```
delete rate-limit <name>
```

Verifying Access Control List Configurations

To verify access control list settings, you can view the access list configuration.

To view the access list configuration use the following command:

```
show access-list {<name> | port <portlist>}
```

To view the rate limit configuration use the following command:

```
show rate-limit {<name> | ports <portlist>}
```

To view the access mask configuration use the following command:

```
show access-mask {<name>}
```

Access Control List Examples

This section presents three access control list examples:

- Using the permit-establish keyword
- Filtering ICMP packets
- Using a rate limit

Using the Permit-Established Keyword

This example uses an access list that permits TCP sessions (Telnet, FTP, and HTTP) to be established in one direction.

The switch, shown in Figure 18, is configured as follows:

- Two VLANs, NET10 VLAN and NET20 VLAN, are defined.
- The NET10 VLAN is connected to port 2 and the NET20 VLAN is connected to port 10
- The IP addresses for NET10 VLAN is 10.10.10.1/24.
- The IP address for NET20 VLAN is 10.10.20.1/24.
- The workstations are configured using addresses 10.10.10.100 and 10.10.20.100.
- IP Forwarding is enabled.

Figure 18: Permit-established access list example topology

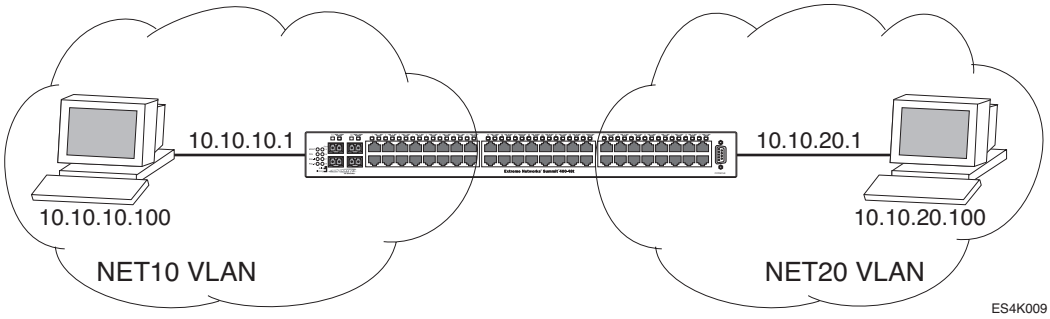

The following sections describe the steps used to configure the example.

Step 1—Deny IP Traffic.

First, create an access-mask that examines the IP protocol field for each packet. Then create two access-lists, one that blocks all TCP, one that blocks UDP. Although ICMP is used in conjunction with IP, it is technically not an IP data packet. Thus, ICMP data traffic, such as ping traffic, is not affected.

The following commands creates the access mask and access lists:

```
create access-mask ipproto_mask ipprotocol ports precedence 25000
create access-list denytcp ipproto_mask ipprotocol tcp ports 2,10 deny
create access-list denyudp ipproto_mask ipprotocol udp ports 2,10 deny
```

Figure 19 illustrates the outcome of the access control list.

Figure 19: Access control list denies all TCP and UDP traffic

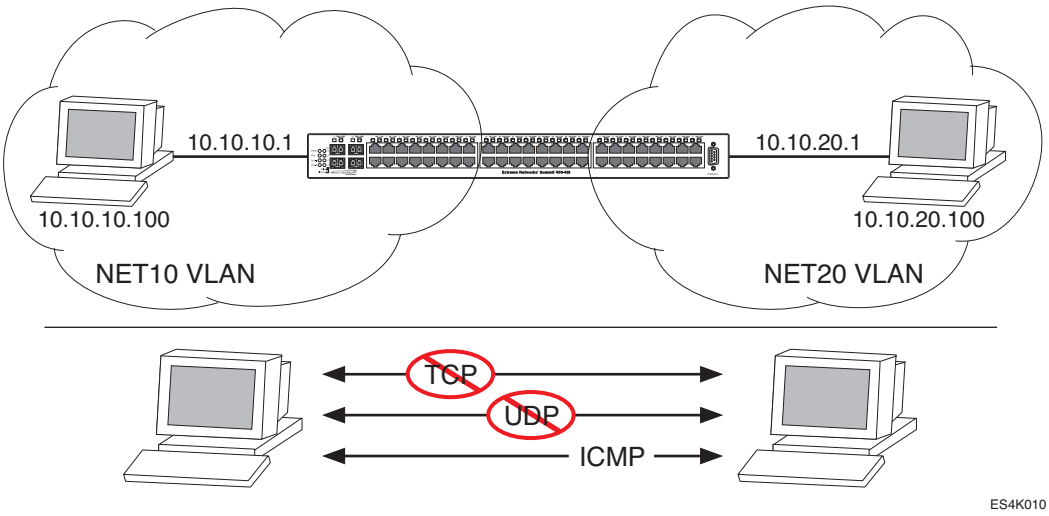

Step 2—Allow TCP traffic.

The next set of access list commands permits TCP-based traffic to flow. Because each session is bi-directional, an access list must be defined for each direction of the traffic flow. UDP traffic is still blocked.

The following commands create the access control list:

```
create access-mask ip_addr_mask ipprotocol dest-ip/32 source-ip/32 ports precedence
20000

create access-list tcp1_2 ip_addr_mask ipprotocol tcp dest-ip 10.10.20.100/32
        source-ip 10.10.10.100/32 ports 2 permit qp1
create access-list tcp2_1 ip_addr_mask ipprotocol tcp dest-ip 10.10.10.100/32
        source-ip 10.10.20.100/32 ports 10 permit qp1
```

Figure 20 illustrates the outcome of this access list.

Figure 20: Access list allows TCP traffic

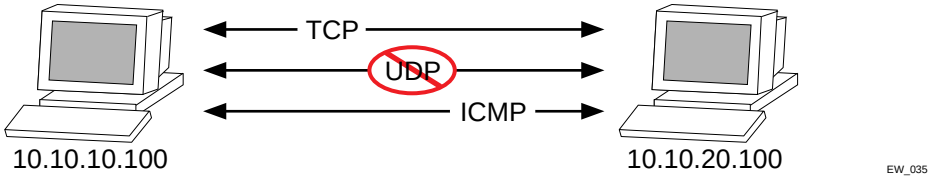

Step 3 - Permit-Established Access List.

When a TCP session begins, there is a three-way handshake that includes a sequence of a SYN, SYN / ACK, and ACK packets. Figure 21 shows an illustration of the handshake that occurs when host A initiates a TCP session to host B. After this sequence, actual data can be passed.

Figure 21: Host A initiates a TCP session to host B

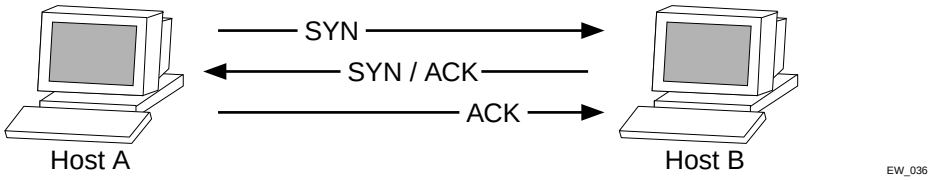

An access list that uses the permit-established keyword filters the SYN packet in one direction.

Use the permit-established keyword to allow only host A to be able to establish a TCP session to host B and to prevent any TCP sessions from being initiated by host B, as illustrated in Figure 21. The commands for this access control list is as follows:

```
create access-mask tcp_connection_mask ipprotocol dest-ip/32 dest-L4port
        permit-established ports precedence 1000
create access-list telnet-deny tcp_connection_mask ipprotocol tcp dest-ip
        10.10.10.100/32 dest-L4port 23 ports 10 permit-established
```

NOTE

This step may not be intuitive. Pay attention to the destination and source address, the ingress port that the rule is applied to, and the desired affect.

NOTE

This rule has a higher precedence than the rule "tcp2_1" and "tcp1_2".

Figure 22 shows the final outcome of this access list.

Figure 22: Permit-established access list filters out SYN packet to destination

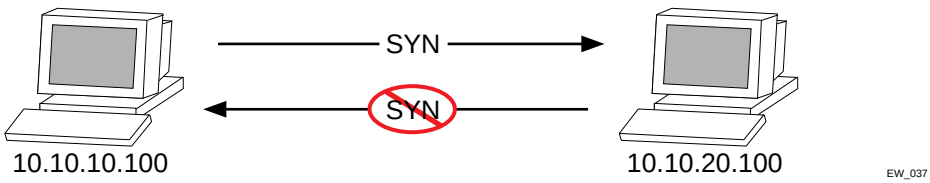

Example 2: Filter ICMP Packets

This example creates an access list that filters out ping (ICMP echo) packets. ICMP echo packets are defined as type 8 code 0.

The commands to create this access control list is as follows:

```
create access-mask icmp_mask ipprotocol icmp-type icmp-code
create access-list denyping icmp_mask ipprotocol icmp icmp-type 8 icmp-code 0 deny
```

The output for this access list is shown in Figure 23.

Figure 23: ICMP packets are filtered out

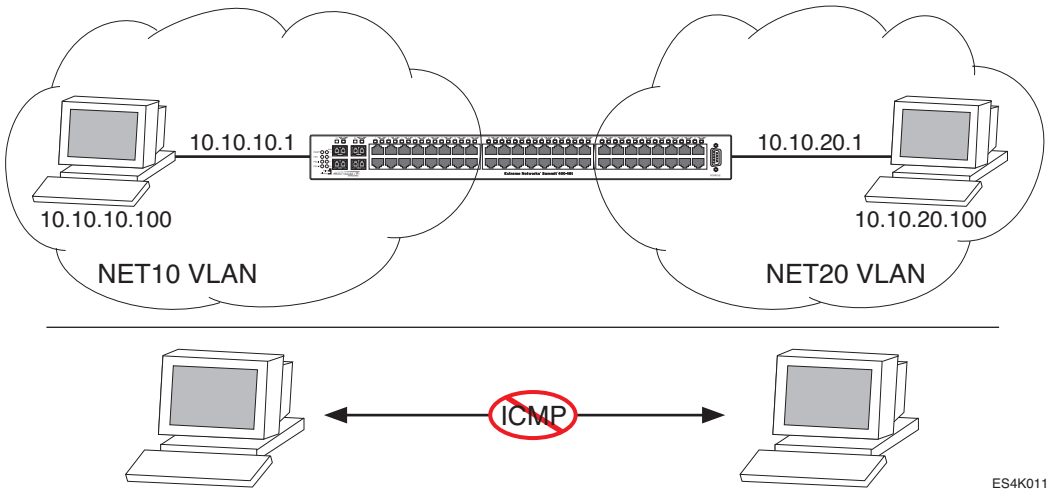

Example 3: Rate-limiting Packets

This example creates a rate limit to limit the incoming traffic from the 10.10.10.x subnet to 10 Mbps on ingress port 2. Ingress traffic on port 2 below the rate limit is sent to QoS profile *qp1* with its DiffServ code point set to 7. Ingress traffic on port 2 in excess of the rate limit will be dropped.

The commands to create this rate limit is as follows:

```
create access-mask port2_mask source-ip/24 ports precedence 100
create rate-limit port2_limit port2_mask source-ip 10.10.10.0/24 port 2 permit qp1 set
code-point 7 limit 10 exceed-action drop
```

Network Login

Network Login is a feature designed to control the admission of user packets into a network by giving addresses only to users that have been properly authenticated. Network Login is controlled by an administrator on a per port, per VLAN basis. When Network Login is enabled on a port in a VLAN, that port will not forward any packets until authentication takes place.

Once Network Login has been enabled on a switch port, that port is placed in a non-forwarding state until authentication takes place. To authenticate, a user (supplicant) must provide the appropriate credentials. These credentials are either approved, in which case the port is placed in forwarding mode, or not approved, and the port remains blocked. The user logout can be initiated by FDB aging or by submitting a logout request.

There are two types of authentication available to use with Network Login: web-based or 802.1x. There are also two different modes of operation available to use with Network Login: Campus mode and ISP mode. The authentication types and modes of operation can be used in any combination. The following sections describe these choices.

Authentication Types

Authentication is handled either as a web-based process or as described in the IEEE 802.1x specification. The initial release of Network Login by Extreme Networks supported only web-based authentication, but later releases have supported both types of authentication.

Although somewhat similar in design and purpose, web-based and 802.1x authentication of Network Login can be considered complementary, with Extreme Networks offering a smooth transition from web-based to 802.1x authentication. In fact, both web-based and 802.1x can be configured on the same switch port. The switch can play the role of the authentication server and authenticate based on its local database of username and password for web-based authentication; or a RADIUS server can be used as the authentication server for web-based and 802.1x authentication.

802.1x Authentication

802.1x will soon be considered the standard for network access authentication. 802.1x authentication currently requires software installed on the client workstation, making it less suitable for a user walk-up scenario, such as a cyber-café or coffee shop. 802.1x authentication also requires an Extensible Authentication Protocol (EAP) capable RADIUS server.

A workstation running Windows XP supports 802.1x natively, and does not require additional authentication software.

Extreme Networks uses a combination of secure certificates and RADIUS server to authenticate the user and configure the switch so that the user is placed on the correct VLAN. When a new user accesses the network, 802.1x authenticates the user through a RADIUS server to a user in an NT domain. The reply from the RADIUS server checks the groups to which the user belongs and then responds to the switch with the proper VLAN. The user is then able to connect to all the resources of the appropriate group after logging in to the network.

Web-Based Authentication

Web-based Network Login does not require any specific client software and can work with any HTTP compliant web browser.

DHCP is needed for web-based network login because the underlying protocol used to carry authentication request-response is HTTP. The client needs an IP address to send and receive HTTP packets. However, before the client is authenticated, the only connection is to the authenticator itself. As a result, the authenticator must be furnished with a temporary DHCP server to distribute the IP address.

The DHCP allocation for Network Login has short time duration of 10 seconds (default value). It is intended to perform web-based network login only. As soon as the client is authenticated, it is deprived of this address. Then it has to go to some other DHCP server in the network to obtain a permanent address, as is normally done. (DHCP is not required for 802.1x because 802.1x uses only layer-2 frames (EAPOL).)

URL redirection is a web-based mechanism to redirect any HTTP request to the base URL of the authenticator when the port is in unauthenticated mode. In other words when user is trying to login to the network using the browser, it will be first redirected to the Network Login page. Only after a successful login will the user be connected to the network.

Co-existence of Web-Based and 802.1x Authentication

ExtremeWare supports both web-based and 802.1x authentication. Authenticating with 802.1x does not require any additional commands besides those used for web-based mode.

When a port is configured for Network Login, the port is put in unauthenticated state. It is ready to perform either type of authentication. Whether to perform web-based or 802.1x is dependent on the type of packets being received from the client. Web-based mode uses HTTP, while 802.1x uses EAPOL with an Ethertype of 0x888e.

This implementation provides a smooth migration path from non-802.1x clients to 802.1x clients. The advantage of web-based mode is platform-independence. While 802.1x mode is currently supported natively only on Windows XP clients, any device with an Internet browser can perform web-based Network Login.

Comparison of Web-Based and 802.1x Authentication

Pros of 802.1x authentication:

- In cases where the 802.1x is natively supported, login and authentication happens transparently.
- Authentication happens at layer 2. Does not involve getting a temporary IP address and subsequent release of the address to a get a more permanent IP address.
- Allows for periodic, transparent, re-authorization of supplicants.

Cons of 802.1x authentication:

- 802.1x native support is available only on the newer operating systems like Windows XP.
- 802.1x requires an EAP-capable RADIUS server.
- TLS authentication method involves Public Key Infrastructure, which requires more administration.
- TTLS is still a Funk / Certicom IETF draft proposal and not a fully accepted standard, but it is easy to deploy and administer.

Pros of web-based authentication:

- Works with any operating system with a web browser. There is no need for any client side software.
- Provides easier administration based on username and password.

Cons of web-based authentication:

- Login process involves juggling with IP addresses and has to be done outside the scope of a regular computer login, therefore it is not tied to Windows login. One has to specifically bring up a login page and initiate a login.
- Supplicants cannot be re-authenticated transparently. Cannot be re-authenticated from the authenticator side.
- Does not support more secure methods of authentication.

Authentication Methods

The authentication methods supported are a matter between the supplicant and the authentication server. The most commonly used methods are:

- MD5-Challenge.
- Transport Layer Security (TLS), which uses Public Key Infrastructure (PKI) and strong mutual authentication.
- Tunneled TLS (TTLS), which is a Funk / Certicom proposal.

TLS represents the most secure protocol among these methods. TTLS is advertised to be as strong as TLS. Both TLS and TTLS are certificate-based, which requires setting up a PKI that can issue, renew, and revoke certificates. TTLS offers ease of deployment because it requires only server certificates and the client can use the MD5 mode of username / password authentication.

For information on setting up a PKI configuration, refer to the documentation for your particular RADIUS server and 802.1x client, if using 802.1x authentication.

Modes of Operation

Network login has two modes of operation:

- Campus mode
- ISP mode

Campus Mode

Campus mode is meant for mobile users who tend to move from one port to another and connect at various locations in the network. In Campus mode, the authenticated port is moved from a temporary VLAN to a permanent VLAN, which then has access to external network resources. Campus mode requires the use of a RADIUS server as part of the authentication process.

ISP Mode

ISP mode is meant for users who will connect through the same port and VLAN each time, as though the switch functions as an ISP. In ISP mode, the port and VLAN remain constant. Before the supplicant is authenticated, the port is in an unauthenticated state. Once authenticated, the port will forward packets.

User Accounts

You can create two types of user accounts for authenticating Network Login users:

- netlogin-only enabled
- netlogin-only disabled

Netlogin-Only Enabled

A netlogin-only enabled user can only log in using Network Login and cannot access the switch using the same login.

Add the following line to the RADIUS server dictionary file for netlogin-only enabled users:

```
Extreme:Extreme-Netlogin-Only = Enabled
```

Netlogin-Only Disabled

A netlogin-only disabled user can log in using Network Login and can also access the switch using Telnet, SSH, or HTTP.

Add the following line to the RADIUS server dictionary file for netlogin-only disabled users:

```
Extreme:Extreme-Netlogin-Only = Disabled
```

Interoperability Requirements

For Network Login to operate, the user (supplicant) software and the authentication server must support common authentication methods. Not all combinations provide the appropriate functionality.

Supplicant Side

On the client side, currently, the only platform that natively supports 802.1x is Windows XP, which performs MD5 and TLS. Other 802.1x clients are available that support other operating systems and support mixes of authentication methods.

A Windows XP 802.1x supplicant can be authenticated as a computer or as a user. Computer authentication requires a certificate installed in the computer certificate store, and user authentication requires a certificate installed in the individual user's certificate store.

By default, the XP machine performs computer authentication as soon as the computer is powered on, or at link-up when no user is logged into the machine. User authentication is performed at link-up when the user is logged in.

The XP machine can be configured to perform computer authentication at link-up even if the user is logged in.

Any client with a web browser can interoperate using web-based authentication.

Authentication Server Side

The RADIUS server used for authentication has to be EAP-capable. Consider the following when choosing a RADIUS server:

- The types of authentication methods supported on RADIUS, as mentioned above.
- Need to support both EAP and traditional Username-Password authentication. These are used by Network Login and switch console login respectively.
- Need to support Vendor Specific Attributes (VSA). Some important parameters such as `Extreme-Netlogin-Vlan` (destination vlan for port movement after authentication) and `Extreme-NetLogin-only` (authorization for network login only) are brought back as VSAs.

  Table 28 and Table 29 show VSA definitions for both web-based network login and 802.1x network login.

Table 28: VSA definitions for web-based network login

VSA	Attribute Value	Type	Sent-in	Description
Extreme-Netlogin -Vlan	203	String	Access-Accept	Name of destination VLAN (must already exist on switch) after successful authentication.
Extreme-Netlogin -Url	204	String	Access-Accept	Destination web page after successful authentication.
Extreme-Netlogin -Url-Desc	205	String	Access-Accept	Text description of network login URL attribute.
Extreme-Netlogin -Only	206	Integer	Access-Accept	Determines if user can authenticate via other means, such as telnet, console, SSH, or Vista. A value of “1” (enabled) indicates that the user can only authenticate via network login. A value of zero (disabled) indicates that the user can also authenticate via other methods.

Table 29: VSA definitions for 802.1x network login

VSA	Attribute Value	Type	Sent-in	Description
Extreme-Netlogin -Vlan	203	String	Access-Accept	Name of destination VLAN (must already exist on switch) after successful authentication.

NOTE

The Extreme Networks vendor ID is 1916.

Multiple Supplicant Support

An important enhancement over the IEEE 802.1x standard, is that ExtremeWare supports multiple clients (supplicants) to be individually authenticated on the same port. This feature makes it possible for two client stations to be connected to the same port, with one being authenticated and the other not. A port's authentication state is the logical “OR” of the individual MAC's authentication states. In other words, a port is authenticated if any of its connected clients is authenticated. Multiple clients can be connected to a single port of authentication server through a hub or layer-2 switch.

Multiple supplicants are supported in ISP mode for both web-based and 802.1x authentication. Multiple supplicants are not supported in Campus mode. Versions of ExtremeWare previous to version 7.1.0 did not support multiple supplicants.

The choice of web-based versus 802.1x authentication is again on a per-MAC basis. Among multiple clients on the same port, it is possible that some clients use web-based mode to authenticate, and some others use 802.1x.

There are certain restrictions for multiple supplicant support:

- Web-based mode will not support Campus mode for multiple supplicant because once the first MAC gets authenticated, the port is moved to a different VLAN and therefore other unauthenticated clients (which are still in the original VLAN), cannot have layer 3 message transactions with the authentication server.

- Once the first MAC is authenticated, the port is transitioned to the authenticated state and other unauthenticated MACs can listen to all data destined for the first MAC. This could raise some security concerns as unauthenticated MACs can listen to all broadcast and multicast traffic directed to a Network Login-authenticated port.

Exclusions and Limitations

The following are limitations and exclusions for Network Login:

- All unauthenticated MACs will be seeing broadcasts and multicasts sent to the port if even a single MAC is authenticated on that port.
- Network Login must be disabled on a port before that port can be deleted from a VLAN.
- In Campus mode, once the port moves to the destination VLAN, the original VLAN for that port is not displayed.
- A Network Login VLAN port should be an untagged Ethernet port and should not be a part of following protocols:
  - ESRP
  - STP
  - VLAN Aggregation
  - VLAN Translation
- Network Login is not supported for T1, E1, T3, ATM, PoS and MPLS TLS interfaces.
- No Hitless Failover support has been added for Network Login.
- Rate-limiting is not supported on Network Login ports (both web-based and 802.1x).
- Network Login and MAC-limits cannot be used together on the same switch (see “Network Login” on page 146).
- EAP-NAK cannot be used to negotiate 802.1x authentication types.

Configuring Network Login

The following configuration example demonstrates how users can initially log in using web-based authentication, allowing them limited access to the network in order to download the 802.1x client and a certificate. After the client is configured, the user is then able to access the network by using 802.1x. The example illustrates the following configuration steps:

1. Create a VLAN on all edge switches called “temp,” which is the initial VLAN to which users will connect before they are authenticated.
2. Create a VLAN on all edge and core switches called “guest,” which is the VLAN from which users will access the Certificate Authority and be able to download the 802.1x software.

The following example demonstrates the first network login configuration step for a Summit 48si edge switch:

```
create vlan temp
configure temp ipaddress 192.168.1.1/24
configure temp add port 1-48
configure vlan temp dhcp-address-range 192.168.1.11 - 192.168.1.200
configure vlan temp dhcp-options default-gateway 192.168.1.1
enable netlogin port 1-48 vlan temp
```

Note that the 192.168 IP address range can be used on all switches because the user is on the VLAN only long enough to log in to the network. After the login is complete, the user is switched to a permanent VLAN with a real IP address delivered from a real DHCP server.

The following example demonstrates the second network login configuration step for a Summit 48si edge switch, in which the guest VLAN is created:

```
create vlan guest
configure guest ipa 45.100.1.101/16
configure guest tag 100
configure guest add port 49-50 tagged
enable bootprelay
configure bootprelay add 45.100.2.101
```

These commands create the special VLAN called “guest” on the real area of the network. Special configuration is needed on the RADIUS server to place users on to the appropriate VLAN when they log in as guests. By using network login in this way, the user goes from unauthenticated to a guest authentication with limited access to resources.

Note that the 45.100.x.x VLAN does not need to be able to route. Extra authentication can be enabled on the Certificate Authority server to more firmly verify the identity of users. The 45.100.x.x VLAN will have the Certificate Authority located on it as well as an HTTP/FTP server to allow the user to download the needed files.

Once the user has installed the certificate from the Certificate Authority and downloaded the 802.1x client, the user can reconnect to the network using 802.1x without the need to authenticate via a web browser. The authentication is handled using PEAP and certificates. The user will be placed in the VLAN that is appropriate for that user’s group.

Web-Based Authentication User Login Using Campus Mode

When web-based authentication is used in Campus mode, the user will follow these steps:

1. Set up the Windows IP configuration for DHCP.
2. Plug into the port that has network login enabled.
3. Log in to Windows.
4. Release any old IP settings and renew the DHCP lease.

   This is done differently depending on the version of Windows the user is running:

   — **Windows 9x**—use the `winipcfg` tool. Choose the Ethernet adapter that is connected to the port on which network login is enabled. Use the buttons to release the IP configuration and renew the DHCP lease.

   — **Windows NT/2000**—use the `ipconfig` command line utility. Use the command `ipconfig/release` to release the IP configuration and `ipconfig/renew` to get the temporary IP address from the switch. If you have more than one Ethernet adapter, specify the adapter by using a number for the adapter following the ipconfig command. You can find the adapter number using the command `ipconfig/all`.

At this point, the client will have its temporary IP address. In this example, the client should have obtained the an IP address in the range 198.162.32.20 - 198.162.32.80.

NOTE

The idea of explicit release/renew is required to bring the network login client machine in the same subnet as the connected VLAN. In Campus Mode using web-based authentication, this requirement is mandatory after every logout and before login again as the port moves back and forth between the temporary and permanent VLANs. On other hand in ISP Mode, release/renew of IP address is not required, as the network login client machine stays in the same subnet as the network login VLAN. In ISP mode, when the network login client connects for the first time, it has to make sure that the machine IP address is in the same subnet as the VLAN to which it is connected.

5 Bring up the browser and enter any URL as `http://www.123.net` or `http://1.2.3.4` or switch IP address as http://<IP address>/login (where IP address could be either temporary or Permanent VLAN Interface for Campus Mode). URL redirection redirects any URL and IP address to the network login page This is significant where security matters most, as no knowledge of VLAN interfaces is required to be provided to network login users, as they can login using a URL or IP address.

   A page opens with a link for Network Login.

6 Click the Network Login link.

   A dialog box opens requesting a username and password.

7 Enter the username and password configured on the RADIUS server.

   After the user has successfully logged in, the user will be redirected to the URL configured on the RADIUS server.

During the user login process, the following takes place:

- Authentication is done through the RADIUS server.
- After successful authentication, the connection information configured on the RADIUS server is returned to the switch:
  - the permanent VLAN
  - the URL to be redirected to (optional)
  - the URL description (optional)
- The port is moved to the permanent VLAN.

  You can verify this using the `show vlan` command. For more information on the `show vlan` command, see “Displaying VLAN Settings” on page 94.

After a successful login has been achieved, there are several ways that a port can return to a non-authenticated, non-forwarding state:

- The user successfully logs out using the logout web browser window.
- The link from the user to the switch’s port is lost.
- There is no activity on the port for 20 minutes.
- An administrator changes the port state.

NOTE

Because network login is sensitive to state changes during the authentication process, Extreme Networks recommends that you do not log out until the login process is complete. The login process is complete when you receive a permanent address.

DHCP Server on the Switch

A DHCP server with limited configuration capabilities is included in the switch to provide IP addresses to clients. The DHCP server is not supported as a standalone feature. It is used only as part of the Network Login feature.

DHCP is enabled on a per port, per VLAN basis. To enable or disable DHCP on a port in a VLAN, use one of the following commands:

```
enable dhcp ports <portlist> vlan <vlan name>
disable dhcp ports <portlist> vlan <vlan name>
configure vlan <vlan name> netlogin-lease-timer <seconds>
```

The switch responds to DHCP requests for unauthenticated clients when DHCP parameters such as `dhcp-address-range` and `dhcp-options` are configured on the network login VLAN. The switch can also answer DHCP requests after authentication if DHCP is enabled on the specified port. If you want network login clients to obtain DHCP leases from an external DHCP server elsewhere on the network, then do not enable DHCP on the switch ports.

Displaying DHCP Information

To display the DHCP configuration, including the DHCP range, DHCP lease timer, network login lease timer, DHCP-enabled ports, IP address, MAC address, and time assigned to each end device, use the following command:

```
show vlan <vlan name> dhcp-address-allocation
```

Displaying Network Login Settings

To display the network login settings, use the following command:

```
show netlogin {port <portlist> vlan <vlan name>}
```

Disabling Network Login

Network login must be disabled on a port before you can delete a VLAN that contains that port. To disable network login, use the following command:

```
disable netlogin ports <portlist> vlan <vlan name>
```

Additional Configuration Details

This section discusses additional configuration like switch DNS name, default redirect page, session refresh and logout-privilege. URL redirection requires the switch to be assigned a DNS name. The default name is `network-access.net`. Any DNS query coming to the switch to resolve switch DNS name in unauthenticated mode is resolved by the DNS server on the switch in terms of the interface (to which the network login port is connected) IP-address.

To configure the network login base URL, use the following command:

```
configure netlogin base-url <url>
```

Where <url> is the DNS name of the switch. For example, `configure netlogin base-url network-access.net` makes the switch send DNS responses back to the netlogin clients when a DNS query is made for `network-access.net`.

To configure the network login redirect page, use the following command:

```
configure netlogin redirect-page <url>
```

Where `<url>` defines the redirection information for the users once logged in. This redirection information is used only in case the redirection info is missing from RADIUS server. For example, `configure netlogin base-url http://www.extremenetworks.com` redirects all users to this URL after they get logged in.

To enable or disable the network login session refresh, use one of the following commands:

```
enable netlogin session-refresh {<minutes>}
disable netlogin session-refresh
```

Where `<minutes>` ranges from 1 - 255. The default setting is 3 minutes. `enable netlogin session-refresh {<minutes>}` makes the logout window refresh itself at every configured time interval. `Session -refresh` is disabled by default.

To enable or disable network login logout privilege, use one of the following commands:

```
enable netlogin logout-privilege
disable netlogin logout-privilege
```

This command turns the privilege for netlogin users to logout by popping up (or not popping up) the logout window. `Logout-privilege` is enabled by default.

To enable or disable network login, use one of the following commands:

```
enable netlogin [web-based | dot1x]
disable netlogin [web-based |dot1x]
```

By default netlogin is enabled.

To show all network login parameters, use the following command:

```
show netlogin
```

Switch Protection

Switch protection features enhance the robustness of switch performance. In this category are the following features:

- Routing Access Profiles
- Denial of Service Protection

Routing Access Profiles

Routing access profiles are used to control the advertisement or recognition of routing protocols, such as RIP or OSPF. Routing access profiles can be used to ‘hide’ entire networks, or to trust only specific sources for routes or ranges of routes. The capabilities of routing access profiles are specific to the type of routing protocol involved, but are sometimes more efficient and easier to implement than access lists.

Using Routing Access Profiles

To use routing access profiles, you must perform the following steps:

1. Create an access profile.
2. Configure the access profile to be of type *permit*, *deny*, or *none*.
3. Add entries to the access profile. Entries can be one of the following types:
   - IP addresses and subnet masks
   - VLAN
4. Apply the access profile.

Creating an Access Profile

The first thing to do when using routing access profiles is to create an *access profile*. An access profile has a unique name and contains one of the following entry types:

- A list of IP addresses and associated subnet masks
- A VLAN

You must give the access profile a unique name (in the same manner as naming a VLAN, protocol filter, or Spanning Tree Domain). To create an access profile, use the following command:

```
create access-profile <access profile> type [ipaddress | ipx-node | ipx-net |
ipx-sap | as-path]
```

Configuring an Access Profile Mode

After the access profile is created, you must configure the access profile mode. The access profile mode determines whether the items in the list are to be permitted access or denied access.

Three modes are available:

- **Permit**—The permit access profile mode permits the operation, as long as it matches any entry in the access profile. If the operation does not match any entries in the list, the operation is denied.
- **Deny**—The deny access profile mode denies the operation, as long as it matches any entry in the access profile. If it does not match all specified entries in the list, the operation is permitted.
- **None**—Using the none mode, the access profile can contain a combination of permit and deny entries. Each entry must have a permit or deny attribute. The operation is compared with each entry in the list. Once a match is found, the operation is either permitted or denied, depending on the configuration of the matched entry. If no match is found, the operation is implicitly denied.

To configure the access profile mode, use the following command:

```
configure access-profile <access profile> mode [permit | deny | none]
```

Adding an Access Profile Entry

Next, configure the access profile, using the following command:

```
configure access-profile <access profile> add {<seq_number>} {permit | deny}
[ipaddress <ip address> <mask> {exact} | as-path <path-expression> | bgp-community
[internet | no-export | no-advertise | no-export-subconfed | <as_no:number> | number
<community>] | ipxnet <netid> <netid mask> | ipxsap <sap_type> <service_name> | vlan]
```

The following sections describe the `configure access-profile add` command.

Specifying Subnet Masks

The subnet mask specified in the access profile command is interpreted as a *reverse mask*. A reverse mask indicates the bits that are significant in the IP address. In other words, a reverse mask specifies the part of the address that must match the IP address to which the profile is applied.

If you configure an IP address that is an exact match that is specifically denied or permitted, use a mask of /32 (for example, 141.251.24.28/32). If the IP address represents all addresses in a subnet address that you want to deny or permit, then configure the mask to cover only the subnet portion (for example, 141.251.10.0/24). The keyword `exact` can be used when you wish to match only against the subnet address, and ignore all addresses within the subnet.

If you are using off-byte boundary subnet masking, the same logic applies, but the configuration is more tricky. For example, the network address 141.251.24.128/27 represents any host from subnet 141.251.24.128.

Sequence Numbering

You can specify the sequence number for each access profile entry. If you do not specify a sequence number, entries are sequenced in the order they are added. Each entry is assigned a value of 5 more than the sequence number of the last entry.

Permit and Deny Entries

If you have configured the access profile mode to be none, you must specify each entry type as either ‘permit’ or ‘deny’. If you do not specify the entry type, it is added as a permit entry. If you have configured the access profile mode to be `permit` or deny, it is not necessary to specify a type for each entry.

Autonomous System Expressions

The `AS-path` keyword uses a regular expression string to match against the AS path. Regular expression notation can include any of the characters listed in Table 30.

Table 30: Regular Expression Notation

Character	Definition
N	As number
N_1 - N_2	Range of AS numbers, where N_1 and N_2 are AS numbers and $N_1 < N_2$
[N_x ... N_y]	Group of AS numbers, where N_x and N_y are AS numbers or a range of AS numbers
[^N_x ... N_y]	Any AS numbers other than the ones in the group

Table 30: Regular Expression Notation (Continued)

Character	Definition
.	Matches any number
^	Matches the beginning of the AS path
$	Matches the end of the AS path
–	Matches the beginning or end, or a space
-	Separates the beginning and end of a range of numbers
*	Matches 0 or more instances
+	Matches 1 or more instances
?	Matches 0 or 1 instance
{	Start of AS SET segment in the AS path
}	End of AS SET segment in the AS path
(	Start of a confederation segment in the AS path
)	End of a confederation segment in the AS path

Autonomous System Expression Example

The following example uses combinations of the autonomous system expressions to create a complicated access profile:

```
create access-profile AS1 type as-path
configure access-profile AS1 mode none
```

These commands create the access profile.

```
configure access-profile AS1 add 5 permit as-path "^65535$"
```

This command configures the access profile to permit AS paths that contain only (begin and end with) AS number 65535.

```
configure access-profile AS1 add 10 permit as-path "^65535 14490$"
```

This command configures the access profile to permit AS paths beginning with AS number 65535, ending with AS number 14490, and containing no other AS paths.

```
configure access-profile AS1 add 15 permit as-path "^1 2-8 [11 13 15]$"
```

This command configures the access profile to permit AS paths beginning with AS number 1, followed by any AS number from 2 - 8, and ending with either AS number 11, 13, or 15.

```
configure access-profile AS1 add 20 deny as-path "111 [2-8]"
```

This command configures the access profile to deny AS paths beginning with AS number 111 and ending with any AS number from 2 - 8.

```
configure access-profile AS1 add 25 permit as-path "111 .?"
```

This command configures the access profile to permit AS paths beginning with AS number 111 and ending with any additional AS number, or beginning and ending with AS number 111.

Deleting an Access Profile Entry

To delete an access profile entry, use the following command:

```
configure access-profile <access profile> delete <seq_number>
```

Applying Access Profiles

Once the access profile is defined, apply it to one or more routing protocols or VLANs. When an access profile is applied to a protocol function (for example, the export of RIP routes) or a VLAN, this forms an access policy. A profile can be used by multiple routing protocol functions or VLANs, but a protocol function or VLAN can use only one access profile.

Routing Profiles for RIP

If you are using the RIP protocol, the switch can be configured to use an access profile to determine:

- **Trusted Neighbor**—Use an access profile to determine trusted RIP router neighbors for the VLAN on the switch running RIP. To configure a trusted neighbor policy, use the following command:

  ```
  configure rip vlan [<vlan name> | all] trusted-gateway [<access profile> | none]
  ```

- **Import Filter**—Use an access profile to determine which RIP routes are accepted as valid routes. This policy can be combined with the trusted neighbor policy to accept selected routes only from a set of trusted neighbors. To configure an import filter policy, use the following command:

  ```
  configure rip vlan [<vlan name> | all] import-filter [<access profile> | none]
  ```

- **Export Filter**—Use an access profile to determine which RIP routes are advertised into a particular VLAN, using the following command:

  ```
  configure rip vlan [<vlan name> | all] export-filter [<access profile> | none]
  ```

Examples

In the example shown in Figure 24, a switch is configured with two VLANs, *Engsvrs* and *Backbone*. The RIP protocol is used to communicate with other routers on the network. The administrator wants to allow all internal access to the VLANs on the switch, but no access to the router that connects to the Internet. The remote router that connects to the Internet has a local interface connected to the corporate backbone. The IP address of the local interface connected to the corporate backbone is 10.0.0.10/24.

Figure 24: RIP access policy example

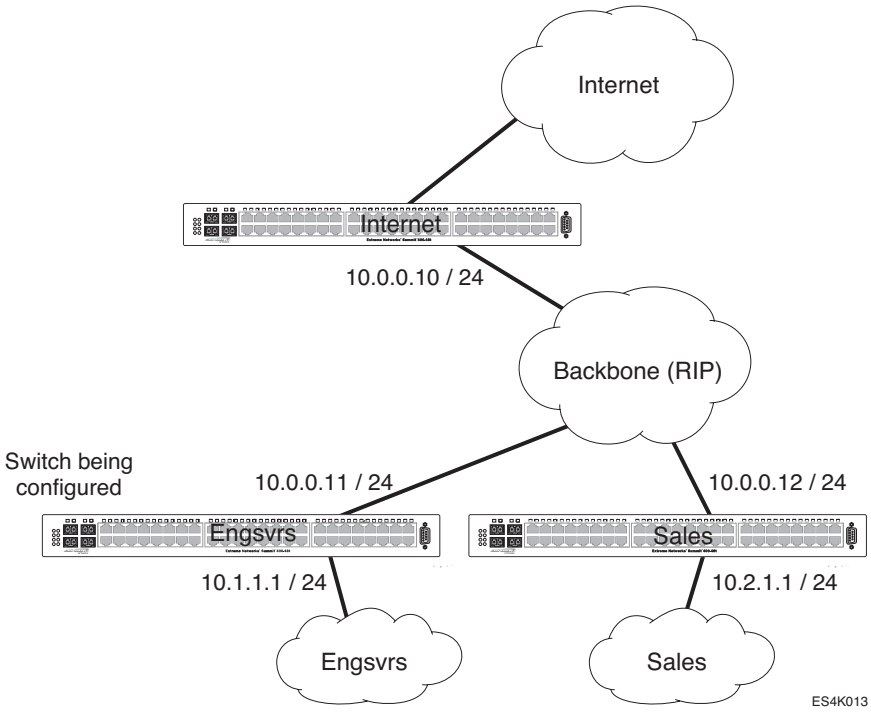

Assuming the backbone VLAN interconnects all the routers in the company (and, therefore, the Internet router does not have the best routes for other local subnets), the commands to build the access policy for the switch would be:

```
create access-profile nointernet ipaddress
configure access-profile nointernet mode deny
configure access-profile nointernet add 10.0.0.10/32
configure rip vlan backbone trusted-gateway nointernet
```

In addition, if the administrator wants to restrict any user belonging to the VLAN *Engsvrs* from reaching the VLAN *Sales* (IP address 10.2.1.0/24), the additional access policy commands to build the access policy would be:

```
create access-profile nosales ipaddress
configure access-profile nosales mode deny
configure access-profile nosales add 10.2.1.0/24
configure rip vlan backbone import-filter nosales
```

This configuration results in the switch having no route back to the VLAN *Sales.*

Routing Access Profiles for OSPF

Because OSPF is a link-state protocol, the access profiles associated with OSPF are different in nature than those associated with RIP. Access profiles for OSPF are intended to extend the existing filtering and security capabilities of OSPF (for example, link authentication and the use of IP address ranges). If you are using the OSPF protocol, the switch can be configured to use an access profile to determine any of the following:

- **Inter-area Filter**—For switches configured to support multiple OSPF areas (an ABR function), an access profile can be applied to an OSPF area that filters a set of OSPF inter-area routes from being sourced from any other areas. To configure an inter-area filter policy, use the following command:

  ```
  configure ospf area <area identifier> interarea-filter [<access profile> | none]
  ```

- **External Filter**—For switches configured to support multiple OSPF areas (an ABR function), an access profile can be applied to an OSPF area that filters a set of OSPF external routes from being advertised into that area. To configure an external filter policy, use the following command:

  ```
  configure ospf area <area identifier> external-filter [<access profile> |none]
  ```

NOTE

If any of the external routes specified in the filter have already been advertised, those routes will remain until the associated LSAs in that area time-out.

- **ASBR Filter**—For switches configured to support RIP and static route re-distribution into OSPF, an access profile can be used to limit the routes that are advertised into OSPF for the switch as a whole. To configure an ASBR filter policy, use the following command:

  ```
  configure ospf asbr-filter [<access profile> | none]
  ```

- **Direct Filter**—For switches configured to support direct route re-distribution into OSPF, an access profile can be used to limit the routes that are advertised into OSPF for the switch as a whole. To configure a direct filter policy, use the following command:

  ```
  configure ospf direct-filter [<access profile> | none]
  ```

Example

Figure 25 illustrates an OSPF network that is similar to the network used previously in the RIP example. In this example, access to the Internet is accomplished by using the ASBR function on the switch labeled Internet. As a result, all routes to the Internet will be done through external routes. Suppose the network administrator wishes to only allow access to certain internet addresses falling within the range 192.1.1.0/24 to the internal backbone.

Figure 25: OSPF access policy example

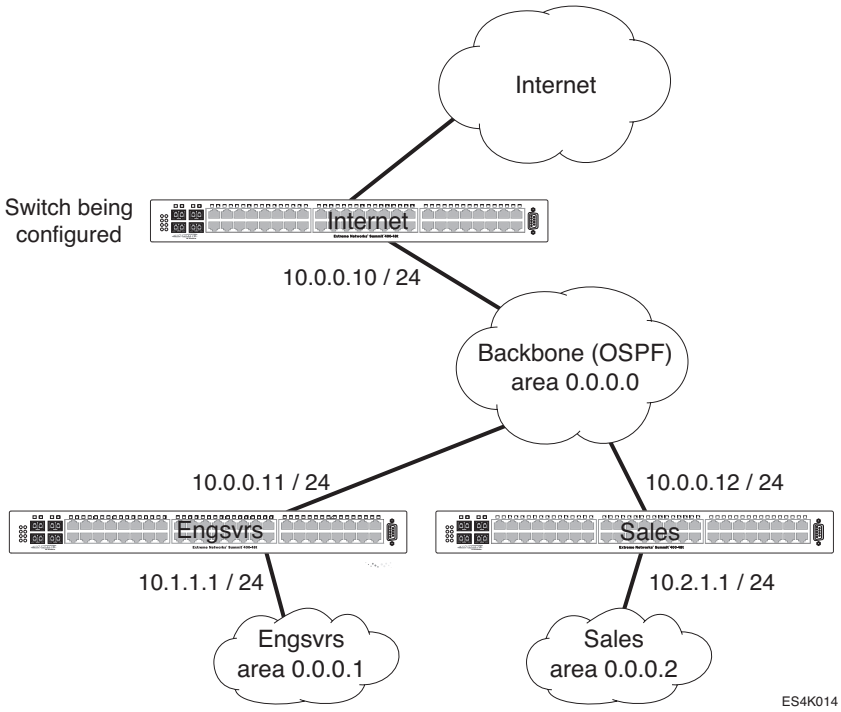

To configure the switch labeled Internet, the commands would be as follows:

```
create access-profile okinternet ipaddress
configure access-profile okinternet mode permit
configure access-profile okinternet add 192.1.1.0/24
configure ospf asbr-filter okinternet
```

Routing Access Profiles for PIM

Because PIM leverages the unicast routing capability that is already present in the switch, the access policy capabilities are, by nature, different. If you are using the PIM protocol for routing IP multicast traffic, you can configure the switch to use an access profile to determine:

Trusted Neighbor—Use an access profile to determine trusted PIM router neighbors for the VLAN on the switch running PIM. To configure a trusted neighbor policy, use the following command:

```
configure pim vlan [<vlan name> | all] trusted-gateway [<access profile> | none]
```

Example

Using PIM, the unicast access profiles can be used to restrict multicast traffic. In this example, a network similar to the example used in the previous RIP example is also running PIM. The network administrator wants to disallow Internet access for multicast traffic to users on the VLAN *Engsvrs*. This is accomplished by preventing the learning of routes that originate from the switch labeled Internet by way of PIM on the switch labeled Engsvrs.

Denial of Service Protection

A Denial-of-Service (DoS) attack occurs when a critical network or computing resource is overwhelmed and rendered inoperative in a way that legitimate requests for service cannot succeed. In its simplest form, a Denial of Service attack is indistinguishable from normal heavy traffic. The Summit 400 switch is not vulnerable to this simple attack because it is designed to process packets in hardware at wire speed. However, there are some operations in any switch or router that are more costly than others, and although normal traffic is not a problem, exception traffic must be handled by the switch's CPU in software.

Some packets that the switch processes in the CPU software include:

- Learning new traffic
- Routing and control protocols including ICMP and OSPF
- Switch management traffic (switch access by Telnet, SSH, HTTP, SNMP, etc...)
- Other packets directed to the switch that must be discarded by the CPU

If any one of these functions is overwhelmed, the CPU can be too busy to service other functions and cause switch performance to suffer. Even with the fast CPU of the Summit 400, there are ways to overwhelm the CPU with packets requiring costly processing.

DoS Protection is designed to help prevent this degraded performance by attempting to characterize the problem and filter out the offending traffic so that other functions can continue. It is the responsibility of DoS Protection to count packets when the switch receives a flood of packets. If the count reaches the threshold, then the flow of these packets to the CPU is blocked.

Configuring Denial of Service Protection

DoS Protection is not enabled on the Summit 400 as a default. To start protecting the switch from attack, first determine what ports are at risk and set limits for the traffic on those ports. Use the following command to identify those ports and to configure the alert-threshold, also known as the disable threshold:

```
configure cpu-dos-protect [ports <portnumber> |all] alert-threshold threshold <pkts>
interval-time <seconds>
```

You can also configure all the ports on the switch to globally implement DoS using the following default values:

- alert-threshold—150 packets per second
- interval-time—1 seconds

To enable all ports on the switch to use DoS Protection, use the following command:

```
enable cpu-dos-protect
```

After enabling DoS Protection, you can use monitor the traffic for the port or the switch by issuing the following command:

```
show cpu-dos-protect [ports <portnumber>]
```

CPU DoS Protection must be enabled for the `show` command to have valid values.

For example, to review the DoS traffic for port 1, issue this command:

```
sh cpu-dos-protect ports 1
```

The output from this command follows:

```
* ex160:22 # sh cpu-dos-protect ports 1

Cpu dos protect: enabled

Port L3Miss  L3Err Bcast IpUnkMcast Learn Curr Int Cfg Thr Cfg Int Pass
_____________________________________________________________________
1    150     150     150  150        150     1       150        1      3

Trusted ports: none
```

The output of this show command displays the following information, which can help you analyze the type of activity coming across the port to the CPU:

- The status of DoS Protection on the port
- Layer 3 miss to the CPU

  These are packets that do not have corresponding IPFDB entries on VLANs, which are enabled for IP forwarding. Packets that are unicasted to the CPU IP are also considered in this category.
- Layer 3 error

  These are IP packets with options, IPMC packets (but not class D address) with checksum errors, and non-IP packets.
- Broadcast traffic
- IP multicast unknown

  These are IPMC packets that do not have corresponding IPMC FDB entries.
- Learning packets

  These are packets that do not have a corresponding FDB entries.
- Current interval

  The current time interval, less than or equal to the configured interval.
- Configured alert threshold

  The maximum number of packets that can be sent to the CPU during the configured interval. This variable is equal to the *configured interval* parameter in seconds for each traffic category.
- Configured interval

  This variable is equal to the *configured interval* parameter in seconds for each traffic category.
- Free pass indicator (Zero in this field indicates a free pass for three intervals after the port comes up.)
- Trusted port status

Creating Trusted Ports

In some cases, traffic from a switch port or group of ports will never cause an attack. These ports can be configured as trusted ports and are not examined under DoS criteria. Trusted ports can prevent innocent hosts from being blocked, or ensure that when an innocent host responds to an attack that the

flood of response packets is not mistaken as the attack. To configure a trusted port, use the following command:

```
configure cpu-dos-protect trusted-ports <port number>
```

For example, to make ports 5 and 7 trusted ports, you would issue this command:

```
config cpu-dos-protect trusted-ports 5, 7
```

To make all ports trusted, or in other words, to disable DoS protection, use the following command:

```
disable cpu-dos-protect
```

Management Access Security

Management access security features control access to the management functions available on the switch. These features help insure that any configuration changes to the switch can only be done by authorized users. In this category are the following features:

- Authenticating Users Using RADIUS or TACACS+
- Secure Shell 2 (SSH2)

Authenticating Users Using RADIUS or TACACS+

ExtremeWare provides two methods to authenticate users who login to the switch:

- RADIUS client
- TACACS+

RADIUS Client

Remote Authentication Dial In User Service (RADIUS, RFC 2138) is a mechanism for authenticating and centrally administrating access to network nodes. The ExtremeWare RADIUS client implementation allows authentication for Telnet, Vista, or console access to the switch.

NOTE

You cannot configure RADIUS and TACACS+ at the same time.

You can define a primary and secondary RADIUS server for the switch to contact. When a user attempts to login using Telnet, http, or the console, the request is relayed to the primary RADIUS server, and then to the secondary RADIUS server, if the primary does not respond. If the RADIUS client is enabled, but access to the RADIUS primary and secondary server fails, the switch uses its local database for authentication.

The privileges assigned to the user (admin versus nonadmin) at the RADIUS server take precedence over the configuration in the local switch database.

To configure the RADIUS servers, use the following command:

```
configure radius [primary | secondary] server [<ipaddress> | <hostname>] {<udp_port>}
client-ip [<ipaddress>]
```

To configure the timeout if a server fails to respond, use the following command:

```
configure radius timeout <seconds>
```

Configuring the Shared Secret Password

In addition to specifying the RADIUS server IP information, RADIUS also contains a means to verify communication between network devices and the server. The *shared secret* is a password configured on the network device and RADIUS server, used by each to verify communication.

To configure the shared secret for RADIUS servers, use the following command:

```
configure radius [primary | secondary] shared-secret {encrypted} [<string>]
```

Enabling and Disabling RADIUS

After server information is entered, you can start and stop RADIUS authentication as many times as necessary without needing to reconfigure server information.

To enable RADIUS authentication, use the following command:

```
enable radius
```

To disable RADIUS authentication, use the following command:

```
disable radius
```

Configuring RADIUS Accounting

Extreme switches are capable of sending RADIUS accounting information. As with RADIUS authentication, you can specify two servers for receipt of accounting information. You can configure RADIUS accounting servers to be the same as the authentication servers, but this is not required.

To specify RADIUS accounting servers, use the following command:

```
configure radius-accounting [primary | secondary] server [<ipaddress> | <hostname>]
{<udp_port>} client-ip [<ipaddress>]
```

To configure the timeout if a server fails to respond, use the following command:

```
configure radius-accounting timeout <seconds>
```

RADIUS accounting also makes use of the shared secret password mechanism to validate communication between network access devices and RADIUS accounting servers.

To specify shared secret passwords for RADIUS accounting servers, use the following command:

```
configure radius-accounting [primary | secondary] shared-secret {encrypted} [<string>]
```

After you configure RADIUS accounting server information, you must enable accounting before the switch begins transmitting the information. You must enable RADIUS authentication for accounting information to be generated. You can enable and disable accounting without affecting the current state of RADIUS authentication.

To enable RADIUS accounting, use the following command:

```
enable radius-accounting
```

To disable RADIUS accounting, use the following command:

```
disable radius-accounting
```

Per-Command Authentication Using RADIUS

The RADIUS implementation can be used to perform per-command authentication. Per-command authentication allows you to define several levels of user capabilities by controlling the permitted command sets based on the RADIUS username and password. You do not need to configure any additional switch parameters to take advantage of this capability. The RADIUS server implementation automatically negotiates the per-command authentication capability with the switch. For examples on per-command RADIUS configurations, see the next section.

Configuring RADIUS Client

You can define primary and secondary server communication information, and for each RADIUS server, the RADIUS port number to use when talking to the RADIUS server. The default port value is 1645. The client IP address is the IP address used by the RADIUS server for communicating back to the switch.

RADIUS RFC 2138 Attributes

The RADIUS RFC 2138 optional attributes supported are as follows:

- User-Name
- User-Password
- Service-Type
- Login-IP-Host

Using RADIUS Servers with Extreme Switches

Extreme Networks switches have two levels of user privilege:

- Read-only
- Read-write

Because there are no CLI commands available to modify the privilege level, access rights are determined when you log in. For a RADIUS server to identify the administrative privileges of a user, Extreme switches expect a RADIUS server to transmit the Service-Type attribute in the Access-Accept packet, after successfully authenticating the user.

Extreme switches grant a RADIUS-authenticated user read-write privilege if a Service-Type value of 6 is transmitted as part of the Access-Accept message from the Radius server. Other Service-Type values, or no value, result in the switch granting read-only access to the user. Different implementations of RADIUS handle attribute transmission differently. You should consult the documentation for your specific implementation of RADIUS when you configure users for read-write access.

Cistron RADIUS

Cistron RADIUS is a popular server, distributed under GPL. Cistron RADIUS can be found at:

http://www.miquels.cistron.nl/radius/

When you configure the Cistron server for use with Extreme switches, you must pay close attention to the users file setup. The Cistron RADIUS dictionary associates the word Administrative-User with Service-Type value 6, and expects the Service-Type entry to appear alone on one line with a leading tab character.

The following is a user file example for read-write access:

```
adminuser   Auth-Type = System

            Service-Type = Administrative-User,

            Filter-Id = "unlim"
```

Livingston (Lucent) RADIUS

Livingston RADIUS is produced by Lucent Technologies primarily for use with their portmaster products. Version 2.1 is released under a BSD license agreement and can be found at ftp://ftp.livingston.com/pub/le/radius/radius21.tar.Z. As with Cistron RADIUS, the Livingston server default dictionary associates Administrative-User with Service-Type value 6. The administrative users file entry example for Cistron RADIUS also works with Livingston RADIUS.

RSA Ace

For users of their SecureID product, RSA offers RADIUS capability as part of their ACE server software. With some versions of ACE, the RADIUS shared-secret is incorrectly sent to the switch resulting in an inability to authenticate. As a work around, do *not* configure a shared-secret for RADIUS accounting and authentication servers on the switch.

Limiting Max-Concurrent Sessions with Funk Software's Steel Belted Radius

For users who have Funk Software's Steel Belted Radius (SBR) server, it is possible to limit the number of concurrent login sessions using the same user account. This feature allows the use of shared user accounts, but limits the number of simultaneous logins to a defined value. Using this feature requires Funk Software Steel-Belted-Radius for Radius Authentication & Accounting.

Complete the following two steps to limit the maximum concurrent login sessions under the same user account:

1. Configure Radius and Radius-Accounting on the switch

   The Radius and Radius-Accounting servers used for this feature must reside on the same physical Radius server. Standard Radius and Radius-Accounting configuration is required as described earlier in this chapter.

2. Modify the Funk SBR 'vendor.ini' file and user accounts

   To configure the Funk SBR server, the file '***vendor.ini***' must be modified to change the Extreme Networks configuration value of '***ignore-ports***' to yes as shown in the example below:

   ```
   vendor-product       = Extreme Networks
   dictionary           = Extreme
   ignore-ports         = yes
   port-number-usage    = per-port-type
   help-id              = 2000
   ```

After modifying the ‘vendor.ini’ file, the desired user accounts must be configured for the Max-Concurrent connections. Using the SBR Administrator application, enable the check box for ‘Max-Concurrent connections’ and fill in the desired number of maximum sessions.

Extreme RADIUS

Extreme Networks provides its users, free of charge, a radius server based on Merit RADIUS. Extreme RADIUS provides per-command authentication capabilities in addition to the standard set of radius features. Source code for Extreme RADIUS can be obtained from the Extreme Networks Technical Assistance Center and has been tested on Red Hat Linux and Solaris.

When Extreme RADIUS is up and running, the two most commonly changed files will be users and profiles. The users file contains entries specifying login names and the profiles used for per-command authentication after they have logged in. Sending a HUP signal to the RADIUS process is sufficient to get changes in the users file to take place. Extreme RADIUS uses the file named profiles to specify command lists that are either permitted or denied to a user based on their login identity. Changes to the profiles file require the RADIUS server to be shutdown and restarted. Sending a HUP signal to the RADIUS process is not enough to force changes to the profiles file to take effect.

When you create command profiles, you can use an asterisk to indicate any possible ending to any particular command. The asterisk cannot be used as the beginning of a command. Reserved words for commands are matched exactly to those in the profiles file. Due to the exact match, it is not enough to simply enter “sh” for “show” in the profiles file, the complete word must be used. Commands can still be entered in the switch in partial format.

When you use per-command authentication, you must ensure that communication between the switch(es) and radius server(s) is not lost. If the RADIUS server crashes while users are logged in, they will have full administrative access to the switch until they log out. Using two RADIUS servers and enabling idle timeouts on all switches will greatly reduce the chance of a user gaining elevated access due to RADIUS server problems.

RADIUS Server Configuration Example (Merit)

Many implementations of RADIUS server use the publicly available Merit© AAA server application, available on the World Wide Web at:

http://www.merit.edu/aaa

Included below are excerpts from relevant portions of a sample Merit RADIUS server implementation. The example shows excerpts from the client and user configuration files. The client configuration file (`ClientCfg.txt`) defines the authorized source machine, source name, and access level. The user configuration file (`users`) defines username, password, and service type information.

`ClientCfg.txt`

```
#Client Name        Key             [type]          [version]   [prefix]
#----------------   --------------  -------------   ---------   --------
#10.1.2.3:256       test             type = nas      v2          pfx
#pm1                %^$%#*(&!(*&)+   type=nas                    pm1.
#pm2                :-):-(;^):-}!    type nas                    pm2.
#merit.edu/homeless hmoemreilte.ses
#homeless           testing          type proxy      v1
#xyz.merit.edu      moretesting      type=Ascend:NAS v1
#anyoldthing:1234   whoknows?        type=NAS+RAD_RFC+ACCT_RFC
10.202.1.3          andrew-linux     type=nas
10.203.1.41         eric             type=nas
```

```
10.203.1.42         eric             type=nas
10.0.52.14          samf             type=nas

users

user    Password = ""
     Filter-Id = "unlim"
admin   Password = "", Service-Type = Administrative
     Filter-Id = "unlim"

eric    Password = "", Service-Type = Administrative
     Filter-Id = "unlim"

albert     Password = "password", Service-Type = Administrative
     Filter-Id = "unlim"

samuel  Password = "password", Service-Type = Administrative
      Filter-Id = "unlim"
```

RADIUS Per-Command Configuration Example

Building on this example configuration, you can use RADIUS to perform per-command authentication to differentiate user capabilities. To do so, use the Extreme-modified RADIUS Merit software that is available from the Extreme Networks by contacting Extreme Networks technical support. The software is available in compiled format for Solaris™ or Linux™ operating systems, as well as in source code format. For all clients that use RADIUS per-command authentication, you must add the following type to the client file:

```
type:extreme:nas + RAD_RFC + ACCT_RFC
```

Within the `users` configuration file, additional keywords are available for `Profile-Name` and `Extreme-CLI-Authorization`. To use per-command authentication, enable the CLI authorization function and indicate a profile name for that user. If authorization is enabled without specifying a valid profile, the user is unable to perform any commands.

Next, define the desired profiles in an ASCII configuration file called `profiles`. This file contains named profiles of exact or partial strings of CLI commands. A named profile is linked with a user through the `users` file. A profile with the `permit on` keywords allows use of only the listed commands. A profile with the `deny` keyword allows use of all commands *except* the listed commands.

CLI commands can be defined easily in a hierarchal manner by using an asterisk (*) to indicate any possible subsequent entry. The parser performs exact string matches on other text to validate commands. Commands are separated by a comma (,) or newline.

Looking at the following example content in profiles for the profile named `PROFILE1`, which uses the `deny` keyword, the following attributes are associated with the user of this profile:

- Cannot use any command starting with `enable`.
- Cannot issue the `disable ipforwarding` command.
- Cannot issue a `show switch` command.
- Can perform all other commands.

We know from the `users` file that this applies to the users `albert` and `lulu`. We also know that `eric` is able to log in, but is unable to perform any commands, because he has no valid profile assigned.

In PROFILE2, a user associated with this profile can use any enable command, the clear counters command and the show management command, but can perform no other functions on the switch. We also know from the users file that gerald has these capabilities.

The following lists the contents of the file users with support for per-command authentication:

```
user    Password = ""
        Filter-Id = "unlim"

admin   Password = "", Service-Type = Administrative
        Filter-Id = "unlim"

eric    Password = "", Service-Type = Administrative, Profile-Name = ""
        Filter-Id = "unlim"
        Extreme:Extreme-CLI-Authorization = Enabled

albert  Password = "", Service-Type = Administrative, Profile-Name =
"Profile1"
          Filter-Id = "unlim"
          Extreme:Extreme-CLI-Authorization = Enabled

lulu    Password = "", Service-Type = Administrative, Profile-Name =
"Profile1"
          Filter-Id = "unlim"
          Extreme:Extreme-CLI-Authorization = Enabled

gerald    Password = "", Service-Type = Administrative, Profile-Name "Profile2"
          Filter-Id = "unlim"
          Extreme:Extreme-CLI-Authorization = Enabled
```

Contents of the file “profiles”:

```
PROFILE1 deny
{
enable  *, disable ipforwarding
show switch
}

PROFILE2
{
enable *, clear counters
show   management
}

PROFILE3 deny
{
create vlan *, configure iproute *, disable *, show fdb
delete *, configure rip add
}
```

Configuring TACACS+

Terminal Access Controller Access Control System Plus (TACACS+) is a mechanism for providing authentication, authorization, and accounting on a centralized server, similar in function to the RADIUS client. The ExtremeWare version of TACACS+ is used to authenticate prospective users who are

attempting to administer the switch. TACACS+ is used to communicate between the switch and an authentication database.

NOTE

You cannot use RADIUS and TACACS+ at the same time.

You can configure two TACACS+ servers, specifying the primary server address, secondary server address, and UDP port number to be used for TACACS+ sessions.

Secure Shell 2 (SSH2)

Secure Shell 2 (SSH2) is a feature of ExtremeWare that allows you to encrypt Telnet session data between a network administrator using SSH2 client software and the switch, or to send encrypted data from the switch to an SSH2 client on a remote system. Image and configuration files may also be transferred to the switch using the Secure Copy Protocol 2 (SCP2). The ExtremeWare CLI provides a command that enable the switch to function as an SSH2 client, sending commands to a remote system via an SSH2 session. It also provides commands to copy image and configuration files to the switch using the SCP2.

The ExtremeWare SSH2 switch application is based on the Data Fellows™ SSH2 server implementation. It is highly recommended that you use the F-Secure® SSH client products from Data Fellows corporation. These applications are available for most operating systems. For more information, see the Data Fellows website at:

http://www.datafellows.com.

NOTE

SSH2 is compatible with the Data Fellows SSH2 client version 2.0.12 or above. SSH2 is not compatible with SSH1.

The ExtremeWare SSH2 switch application also works with SSH2 client and server (version 2.x or later) from SSH Communication Security, and the free SSH2 and SCP2 implementation (version 2.5 or later) from OpenSSH. The SFTP file transfer protocol is required for file transfer using SCP2.

Enabling SSH2 for Inbound Switch Access

Because SSH2 is currently under U.S. export restrictions, you must first obtain a security-enabled version of the ExtremeWare software from Extreme Networks before you can enable SSH2. The procedure for obtaining a security-enabled version of the ExtremeWare software is described in “Security Licensing” on page 31.

You must enable SSH2 on the switch before you can connect to it using an external SSH2 client. Enabling SSH2 involves two steps:

- Enabling SSH2 access, which may include specifying a list of clients that can access the switch, and specifying a TCP port to be used for communication.

  By default, if you have a security license, SSH2 is enabled using TCP port 22, with no restrictions on client access.

- Generating or specifying an authentication key for the SSH2 session.

To enable SSH2, use the following command:

```
enable ssh2 {access-profile [<access profile> | none]} {port <tcp_port_number>}
```

You can specify a list of predefined clients that are allowed SSH2 access to the switch. To do this, you must create an access profile that contains a list of allowed IP addresses.

You can also specify a TCP port number to be used for SSH2 communication. By default the TCP port number is 22.

The supported ciphers are 3DES-CBC and Blowfish. The supported key exchange is DSA.

An authentication key must be generated before the switch can accept incoming SSH2 sessions. This can be done automatically by the switch, or you can enter a previously generated key. To have the key generated by the switch, use the following command:

```
configure ssh2 key
```

You are prompted to enter information to be used in generating the key. The key generation process takes approximately ten minutes. Once the key has been generated, you should save your configuration to preserve the key.

To use a key that has been previously created, use the following command:

```
configure ssh2 key {pregenerated}
```

You are prompted to enter the pregenerated key.

The key generation process generates the SSH2 private host key. The SSH2 public host key is derived from the private host key, and is automatically transmitted to the SSH2 client at the beginning of an SSH2 session.

Before you initiate a session from an SSH2 client, ensure that the client is configured for any nondefault access list or TCP port information that you have configured on the switch. Once these tasks are accomplished, you may establish an SSH2-encrypted session with the switch. Clients must have a valid user name and password on the switch in order to log into the switch after the SSH2 session has been established.

For additional information on the SSH protocol refer to [FIPS-186] Federal Information Processing Standards Publication (FIPSPUB) 186, Digital Signature Standard, 18 May 1994. This can be download from: ftp://ftp.cs.hut.fi/pub/ssh. General technical information is also available from:

http://www.ssh.fi

Using SCP2 from an External SSH2 Client

In ExtremeWare version 6.2.1 or later, the SCP2 protocol is supported for transferring image and configuration files to the switch from the SSH2 client, and for copying the switch configuration from the switch to an SSH2 client.

CAUTION

You can download a configuration to an Extreme Networks switch using SCP. If you do this, you cannot save this configuration. If you save this configuration and reboot the switch, the configuration will be corrupted.

The user must have administrator-level access to the switch. The switch can be specified by its switch name or IP address.

Configuration or image files stored on the system running the SSH2 client may be named as desired by the user. However, files on the switch have predefined names, as follows:

- `configuration.cfg`—The current configuration
- `incremental.cfg`—The current incremental configuration
- `primary.img`—The primary ExtremeWare image
- `secondary.img`—The secondary ExtremeWare image
- `bootrom.img`—The BootROM image

For example, to copy an image file saved as *image1.xtr* to switch with IP address 10.10.0.5 as the primary image using SCP2, you would enter the following command within your SSH2 session:

```
scp image1.xtr admin@10.20.0.5:primary.img
```

To copy the configuration from the switch and save it in file *config1.save* using SCP, you would enter the following command within your SSH2 session:

```
scp admin@10.10.0.5:configuration.cfg config1.save
```

SSH2 Client Functions on the Switch

In ExtremeWare version 6.2.1 or later, an Extreme Networks switch can function as an SSH2 client. This means you can connect from the switch to a remote device running an SSH2 server, and send commands to that device. You can also use SCP2 to transfer files to and from the remote device.

You do not need to enable SSH2 or generate an authentication key to use the SSH2 and SCP2 commands from the ExtremeWare CLI.

To send commands to a remote system using SSH2, use the following command:

```
ssh2 {cipher [3des | blowfish]} {port <portnum>} {compression [on | off]} {user
<username>} {debug <debug_level>} {<username>@} [<host> | <ipaddress>] {<remote
command>}
```

The remote commands can be any commands acceptable by the remote system. You can specify the login user name as a separate argument, or as part of the user@host specification. If the login user name for the remote system is the same as your user name on the switch, you can omit the username parameter entirely.

To initiate a file copy from a remote system to the switch using SCP2, use the following command:

```
scp2 {cipher [3des | blowfish]} {port <portnum>} {debug <debug_level>} <user>@
[<hostname> | <ipaddress>] :<remote_file> [configuration {incremental} | image
[primary | secondary] | bootrom]
```

To initiate a file copy to a remote system from the switch using SCP2, use the following command:

```
scp2 {cipher [3des | blowfish]} {port <portnum>} {debug <debug_level>} configuration
<user>@ [<hostname> | <ipaddress>]:<remote_file>
```

10 Ethernet Automatic Protection Switching

This chapter describes the use of the Ethernet Automatic Protection Switching (EAPS™) protocol, and includes information on the following topics:

Overview of the EAPS Protocol

The EAPS protocol provides fast protection switching to layer 2 switches interconnected in an Ethernet ring topology, such as a Metropolitan Area Network (MAN) or large campuses (see Figure 26).

EAPS protection switching is similar to what can be achieved with the Spanning Tree Protocol (STP), but offers the advantage of converging in less than a second when a link in the ring breaks.

An Ethernet ring built using EAPS can have resilience comparable to that provided by SONET rings, at a lower cost and with fewer restraints (e.g., ring size). The EAPS technology developed by Extreme Networks to increase the availability and robustness of Ethernet rings is described in *RFC 3619: Extreme Networks' Ethernet Automatic Protection Switching (EAPS) Version 1*.

In order to use EAPS, you must enable EDP on the switch and EAPS ring ports. For more information on EDP, see "Extreme Discovery Protocol" on page 85.

EAPS operates by declaring an EAPS domain on a single ring. Any VLAN that warrants fault protection is configured on all ring ports in the ring, and is then assigned to an EAPS domain. On that ring domain, one switch, or node, is designated the *master* node (see Figure 27), while all other nodes are designated as *transit* nodes.

Figure 26: Gigabit Ethernet fiber EAPS MAN ring

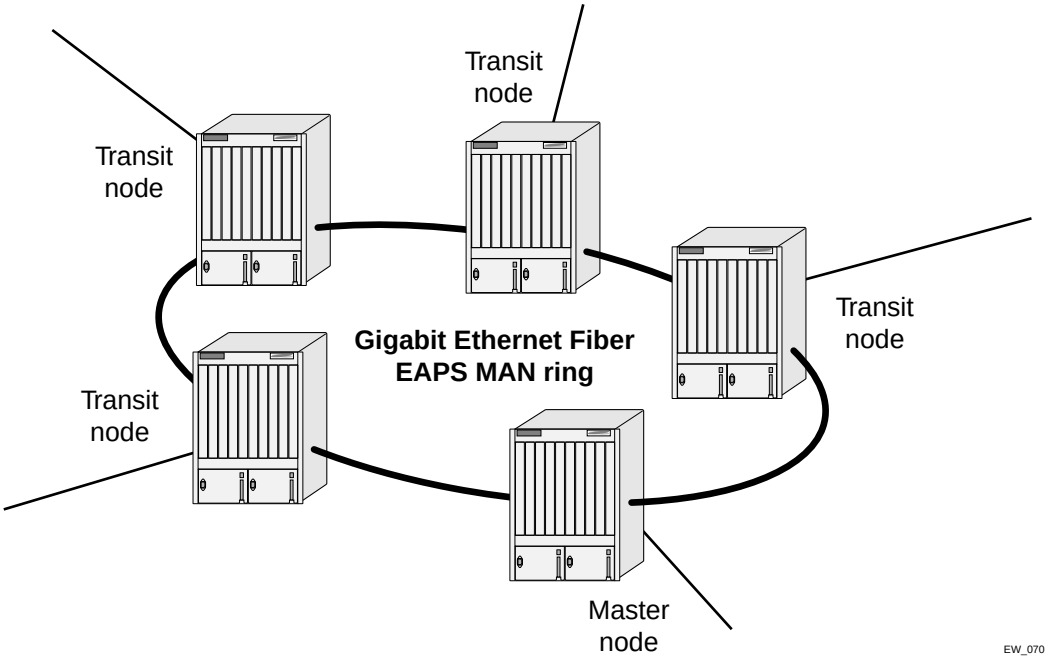

One port of the master node is designated the master node's *primary* port (P) to the ring; another port is designated as the master node's *secondary* port (S) to the ring. In normal operation, the master node blocks the secondary port for all non-control traffic belonging to this EAPS domain, thereby avoiding a loop in the ring, like STP. Layer 2 switching and learning mechanisms operate per existing standards on this ring.

NOTE

Like the master node, each transit node is also configured with a primary port and a secondary port on the ring, but the primary/secondary port distinction is ignored as long as the node is configured as a transit node.

Figure 27: EAPS operation

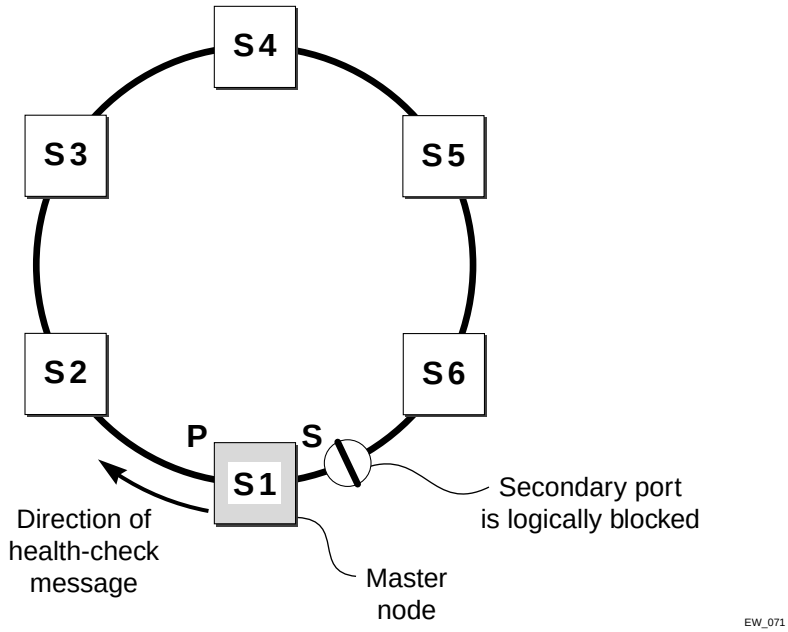

If the ring is complete, the master node logically blocks all data traffic in the transmit and receive directions on the secondary port to prevent a loop. If the master node detects a break in the ring, it unblocks its secondary port and allows data traffic to be transmitted and received through it.

EAPS Terms

Table 31 describes terms associated with EAPS.

Table 31: EAPS Terms

Term	Description
EAPS domain	A domain consists of a series of switches, or nodes, that comprise a single ring in a network. An EAPS domain consists of a master node, transit nodes, and on the master node, one primary port and one secondary port. EAPS operates by declaring an EAPS domain on a single ring.
EDP	Extreme Discovery Protocol. A protocol used to gather information about neighbor Extreme switches. Extreme switches use EDP to exchange topology information.
master node	A switch, or node, that is designated the master in an EAPS domain ring. The master node blocks the secondary port for all non-control traffic belonging to this EAPS domain, thereby avoiding a loop in the ring.
transit node	A switch, or node, that is not designated a master in an EAPS domain ring.
primary port	A port on the master node that is designated the primary port to the ring. The transit node ignores the primary port distinction as long as the node is configured as a transit node.
secondary port	A port on the master node that is designated the secondary port to the ring. The transit node ignores the secondary port distinction as long as the node is configured as a transit node.
control VLAN	A VLAN that sends and receives EAPS messages. You must configure one control VLAN for each EAPS domain.

Table 31: EAPS Terms (Continued)

Term	Description
protected VLAN	A VLAN that carries data traffic through an EAPS domain. You must configure one or more protected VLANs for each EAPS domain. (Also known as data VLAN)

Fault Detection and Recovery

EAPS fault detection on a ring is based on a single *control* VLAN per EAPS domain. This EAPS domain provides protection to one or more data-carrying VLANs called *protected* VLANs.

The control VLAN is used only to send and receive EAPS messages; the protected VLANs carry the actual data traffic. As long as the ring is complete, the EAPS master node blocks the protected VLANs from accessing its secondary port.

NOTE

The control VLAN is not blocked. Messages sent on the control VLAN must be allowed into the switch for the master node to determine whether the ring is complete.

To avoid loops in the network, the control VLAN must be NOT be configured with an IP address, and ONLY ring ports may be added to the VLAN.

Figure 28: EAPS fault detection and protection switching

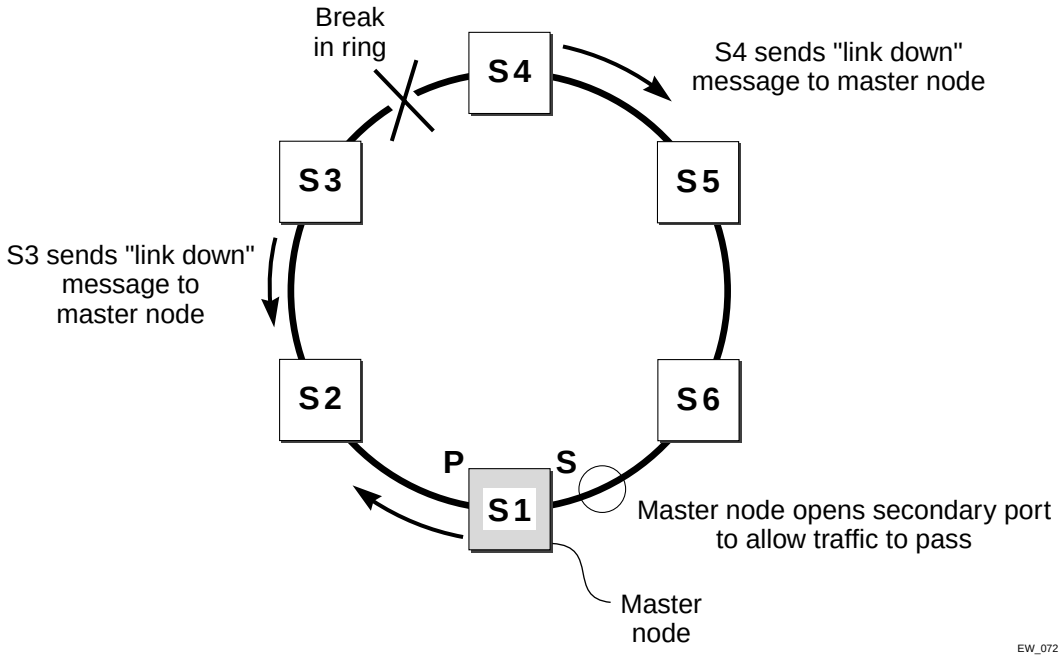

A master node detects a ring fault in one of three ways:

- Link-down message sent by a transit node
- Ring port down event sent by hardware layers

- Polling response

Link Down Message Sent by a Transit Node

When any transit node detects a loss of link connectivity on any of its ring ports, it immediately sends a “link down” message on the control VLAN using its good link to the master node.

When the master node receives the “link down” message (see Figure 28), it immediately declares a “failed” state and opens its logically blocked secondary port on all the protected VLANs. Now, traffic can flow through the master’s secondary port. The master node also flushes its FDB and sends a message on the control VLAN to all of its associated transit nodes to flush their forwarding databases as well, so that all of the switches can learn the new paths to layer 2 end stations on the reconfigured ring topology.

Ring Port Down Event Sent by Hardware Layer

When a ring port goes down on a master node switch, it is notified by the lower hardware layer and immediately goes into a “failed” state.

If the primary ring port goes down, the secondary port is opened. The normal operation of flushing its FDB and sending a “link-down” message to all transit nodes is performed.

Polling

The master node transmits a health-check packet on the control VLAN at a user-configurable interval (see Figure 27). If the ring is complete, the master node will receive the health-check packet on its secondary port (the control VLAN is not blocked on the secondary port). When the master node receives the health-check packet, it resets its failtimer and continues normal operation.

If the master node does not receive the health-check packet before the failtimer interval expires, and the failtime expiry action is set to open the secondary port when the failtimer expires, it declares a “failed” state. The switch then performs the same steps described above: it unblocks its secondary port for access by the protected VLANs, flushes its forwarding database (FDB), and sends a “flush FDB” message to its associated transit nodes.

To change the expiry timer action, use the following command:

```
configure eaps <name> failtime expiry-action [ open-secondary-port | send-alert]
```

To change the duration of the failtime, use the following command:

```
configure eaps <name> failtime [<seconds>]
```

Restoration Operations

The master node continues sending health-check packets out its primary port even when the master node is operating in the failed state. As long as there is a break in the ring, the fail-period timer of the master node will continue to expire and the master node will remain in the failed state.

When the broken link is restored, the master will receive its health-check packet back on its secondary port, and will once again declare the ring to be complete. It will logically block the protected VLANs on its secondary port, flush its FDB, and send a “flush FDB” message to its associated transit nodes.

During the time between when the transit node detects that the link is operable again and when the master node detects that the ring is complete, the secondary port on the master node is still open and data could start traversing the transit node port that just came up. To prevent the possibility of a such a temporary loop, when the transit node detects that its failed link is up again, it will perform these steps:

1. For the port that just came up, put all the protected VLANs traversing that port into a temporary blocked state.
2. Remember which port has been temporarily blocked.
3. Set the state to Preforwarding.

When the master node receives its health-check packet back on its secondary port, and detects that the ring is once again complete, it sends a message to all its associated transit nodes to flush their forwarding databases.

When the transit nodes receive the message to flush their forwarding databases, they perform these steps:

1. Flush their forwarding databases on the protected VLANs.
2. If the port state is set to Preforwarding, unblock all the previously blocked protected VLANs for the port.

Configuring EAPS on a Switch

This section describes how to configure EAPS on a switch.

Creating and Deleting an EAPS Domain

Each EAPS domain is identified by a unique domain name.

To create an EAPS domain, use the following command:

```
create eaps <name>
```

The `name` parameter is a character string of up to 32 characters that identifies the EAPS domain to be created. EAPS domain names and VLAN names must be unique: Do not use the same name string to identify both an EAPS domain and a VLAN.

The following command example creates an EAPS domain named “eaps_1”:

```
create eaps eaps_1
```

To delete an EAPS domain, use the following command:

```
delete eaps <name>
```

The following command example deletes the EAPS domain “eaps_1”:

```
delete eaps eaps_1
```

Defining the EAPS Mode of the Switch

To configure the EAPS node type of the switch, use the following command:

```
configure eaps <name> mode [master | transit]
```

One node on the ring must be configured as the master node for the specified domain; all other nodes on the ring are configured as transit nodes for the same domain.

The following command example identifies this switch as the master node for the EAPS domain named eaps_1.

```
configure eaps eaps_1 mode master
```

The following command example identifies this switch as a transit node for the EAPS domain named eaps_1.

```
configure eaps eaps_1 mode transit
```

Configuring EAPS Polling Timers

To set the values of the polling timers the master node uses for the EAPS health-check packet that is circulated around the ring for an EAPS domain, use the following command:

```
configure eaps <name> hellotime <seconds>
configure eaps <name> failtime [<seconds>]
```

To configure the action taken if there is a break in the ring, use the following command:

```
configure eaps <name> failtime expiry-action [ open-secondary-port | send-alert]
```

NOTE

These commands apply only to the master node. If you configure the polling timers for a transit node, they will be ignored. If you later reconfigure that transit node as the master node, the polling timer values will be used as the current values.

Use the `hellotime` keyword and its associated `seconds` parameter to specify the amount of time the master node waits between transmissions of health-check packets on the control VLAN. `seconds` must be greater than 0 when you are configuring a master node. The default value is one second.

NOTE

Increasing the `hellotime` *value keeps the processor from sending and processing too many health-check packets. Increasing the* `hellotime` *value should not affect the network convergence time, because transit nodes are already sending "link down" notifications.*

Use the `failtime` keyword and `seconds` parameters to specify the amount of time the master node waits before the failtimer expires.

The `seconds` parameter must be greater than the configured value for `hellotime`. The default value is three seconds.

You can configure the action taken when the failtimer expires by using the `configure eaps failtime expiry-action` command. Use the `send-alert` parameter to send an alert when the failtimer expires. Instead of going into a "failed" state, the master node remains in a "Complete" or "Init" state, maintains

the secondary port blocking, and writes a critical error message to syslog warning the user that there is a fault in the ring. An SNMP trap is also sent.

To use the failtimer expiry action of earlier releases, use the `open-secondary-port` parameter.

NOTE

Increasing the failtime value provides more protection by waiting longer to receive a health-check packet when the network is congested.

The following command examples configure the hellotime value for the EAPS domain "eaps_1" to 2 seconds, the failtime value to 15 seconds, and the failtime expiry-action to open the secondary port if the failtimer expires:

```
configure eaps eaps_1 hellotime 2
configure eaps eaps_1 failtime 15
configure eaps eaps_1 failtimer expiry-action open-secondary-port
```

Configuring the Primary and Secondary Ports

Each node on the ring connects to the ring through two ring ports. As part of the protection switching scheme, one port must be configured as the *primary* port; the other must be configured as the *secondary* port.

If the ring is complete, the master node prevents a loop by logically blocking all data traffic in the transmit and receive directions on its secondary port. If the master node subsequently detects a break in the ring, it unblocks its secondary port and allows data traffic to be transmitted and received through it.

To configure a node port as primary or secondary, use the following command:

```
configure eaps <name> [primary | secondary] port <port number>
```

The following command example adds port 1 of the Summit 400-48 switch to the EAPS domain "eaps_1" as the primary port.

```
configure eaps eaps_1 primary port 1
```

Configuring the EAPS Control VLAN

You must configure one *control* VLAN for each EAPS domain. The control VLAN is used only to send and receive EAPS messages.

NOTE

If the domain is active, you cannot delete the domain or modify the configuration of the control VLAN.

To configure the EAPS control VLAN for the domain, use the following command:

```
configure eaps <name> add control vlan <vlan_name>
```

NOTE

The control VLAN must NOT be configured with an IP address. In addition, only ring ports may be added to this control VLAN. No other ports can be members of this VLAN. Failure to observe these restrictions can result in a loop in the network.

NOTE

When you configure the VLAN that will act as the control VLAN, that VLAN must be assigned a QoS profile of Qp8, and the ring ports of the control VLAN must be tagged.

By assigning the control VLAN a QoS profile of Qp8 (with the QoS profile `HighHi` priority setting), you ensure that EAPS control VLAN traffic is serviced before any other traffic and that control VLAN messages reach their intended destinations. For example, if the control VLAN is not assigned the highest priority and a broadcast storm occurs in the network, the control VLAN messages might be dropped at intermediate points. Assigning the control VLAN the highest priority prevents dropped control VLAN messages.

Because the QoS profile `High` priority setting by itself should ensure that the control VLAN traffic gets through a congested port first, you should not need to set the QoS profile minimum bandwidth (`minbw`) or maximum bandwidth (`maxbw`) settings. However, if you plan to use QoS (profile priority and bandwidth settings) for other traffic, you might need to set a `minbw` value on Qp8 for control VLAN traffic. Whether you need to do this depends entirely on your configuration.

The following command example adds the control VLAN “keys” to the EAPS domain “eaps_1”.

```
configure eaps eaps_1 add control vlan keys
```

Configuring the EAPS Protected VLANs

You must configure one or more *protected* VLANs for each EAPS domain. The protected VLANs are the data-carrying VLANs.

NOTE

When you configure the VLAN that will act as a protected VLAN, the ring ports of the protected VLAN must be tagged (except in the case of the default VLAN).

To configure an EAPS protected VLAN, use the following command:

```
configure eaps <name> add protect vlan <vlan_name>
```

NOTE

As long as the ring is complete, the master node blocks the protected VLANs on its secondary port.

The following command example adds the protected VLAN “orchid” to the EAPS domain “eaps_1.”

```
configure eaps eaps_1 add protect vlan orchid
```

NOTE

The configuration of the Superbridge, SubBridge, and IP range control VLANs cannot be modified.

Enabling and Disabling an EAPS Domain

To enable a specific EAPS domain, use the following command:

```
enable eaps {<name>}
```

To disable a specific EAPS domain, use the following command:

```
disable eaps {<name>}
```

Enabling and Disabling EAPS

To enable the EAPS function for the entire switch, use the following command:

```
enable eaps
```

To disable the EAPS function for the entire switch, use the following command:

```
disable eaps
```

Unconfiguring an EAPS Ring Port

Unconfiguring an EAPS port sets its internal configuration state to INVALID, which causes the port to appear in the Idle state with a port status of Unknown when you use the `show eaps {<name>} {detail}` command to display the status information about the port.

To unconfigure an EAPS primary or secondary ring port for an EAPS domain, use the following command:

```
unconfigure eaps <name> [primary | secondary] port
```

The following command example unconfigures this node's EAPS primary ring port on the domain "eaps_1":

```
unconfigure eaps eaps_1 primary port
```

Displaying EAPS Status Information

To display EAPS status information, use the following command:

```
show eaps summary
```

The results for this command are as follows:

```
EAPS Enabled: Yes
Number of EAPS instances: 1
EAPSD-Bridge links: 2

                                    Pri     Sec                          Vlan
Domain        State        Mo En   Port    Port     Control-Vlan (VID) count
------------  ------------ -- -- ------- -------   ------------------ -----

eaps1         Complete     M  Y    10      20      cvlan       (0100)    1
```

To display more detailed EAPS status information, use the following command:

```
show eaps {<name>} {detail}
```

If you enter the `show eaps` command without an argument or keyword, the command displays a summary of status information for all configured EAPS domains. You can use the `detail` keyword to display more detailed status information.

NOTE

The output displayed by this command depends on whether the node is a transit node or a master node. The display for a transit node contains information fields that are not shown for a master node. Also, some state values are different on a transit node than on a master node.

The following example of the `show eaps {<name>} {detail}` command displays detailed EAPS information for a transit node. Table 32 describes the fields and values in the display.

```
EAPS Enabled: Yes
Number of EAPS instances: 1
EAPSD-Bridge links: 2

  Name: "eaps1" (instance=0)
  State: Links-Up        [Running: Yes]
  Enabled: Yes    Mode: Transit
  Primary port: 10             Port status: Up        Tag status:Tagged
  Secondary port: 20           Port status: Up        Tag status:Tagged
  Hello Timer interval: 1 sec     Fail Timer interval: 3 sec
  Preforwarding Timer interval: 6 sec
  Last update: From Master Id 00:04:96:14:46:B0, at Wed Jan 28 15:38:16
2004
  EAPS Domain has following Controller Vlan:
    Vlan Name            VID             QosProfile
    "cvlan"              0100            QP8
  EAPS Domain has following Protected Vlan(s):
    Vlan Name            VID             QosProfile
    "pvlan"              0200            QP1
  Number of Protected Vlans: 1
```

Table 32: show eaps Display Fields

Field	Description
EAPS Enabled:	Current state of EAPS on this switch: • Yes—EAPS is enabled on the switch. • No—EAPS is not enabled.
Number of EAPS instances:	Number of EAPS domains created. The maximum number of EAPS domains per switch is 64.
EAPSD-Bridge links:	The total number of EAPS bridge links in the system. The maximum count is 4096. Each time a VLAN is added to EAPS, this count increments by 1.
Name:	The configured name for this EAPS domain.
(Instance=)	The instance number is created internally by the system.

Table 32: show eaps Display Fields (Continued)

Field	Description
State:	On a transit node, the command displays one of the following states: • Idle—The EAPS domain has been enabled, but the configuration is not complete. • Links-Up—This EAPS domain is running, and both its ports are up and in the FORWARDING state. • Links-Down—This EAPS domain is running, but one or both of its ports are down. • Preforwarding—This EAPS domain is running, and both of its ports are up, but one of them is in a temporary BLOCKED state. On a master node, the command displays one of the following states: • Idle—The EAPS domain has been enabled, but the configuration is not complete. • Init—The EAPS domain has started but has not yet determined the status of the ring. The secondary port is in a BLOCKED state. • Complete—The ring is in the COMPLETE state for this EAPS domain. • Failed—There is a break in the ring for this EAPS domain. • [Failtimer Expired]—When the failtimer expires and it's action is set to send-alert, this flag is set. This flag indicates there is a misconfiguration or hardware problem in the EAPS ring. The EAPS master node will continue to remain in COMPLETE or INIT state with it's secondary port blocking.
[Running: …]	• Yes—This EAPS domain is running. • No—This EAPS domain is not running.
Enabled:	Indicates whether EAPS is enabled on this domain. • Y—EAPS is enabled on this domain. • N—EAPS is not enabled.
Mode:	The configured EAPS mode for this switch: transit (T) or master (M).
Primary/Secondary port:	The port numbers assigned as the EAPS primary and secondary ports. On the master node, the port distinction indicates which port is blocked to avoid a loop.
Port status:	• Unknown—This EAPS domain is not running, so the port status has not yet been determined. • Up—The port is up and is forwarding data. • Down—The port is down. • Blocked—The port is up, but data is blocked from being forwarded.
Tag status:	Tagged status of the control VLAN: • Tagged—The control VLAN has this port assigned to it, and the port is tagged in the VLAN. • Untagged—The control VLAN has this port assigned to it, but the port is untagged in the control VLAN. • Undetermined—Either a VLAN has not been added as the control VLAN to this EAPS domain or this port has not been added to the control VLAN.
Hello Timer interval:	The configured value of the timer in seconds, specifying the time that the master node waits between transmissions of health-check packets.

Table 32: show eaps Display Fields (Continued)

Field	Description
Fail Timer interval:	The configured value of the timer in seconds, specifying the time that the master node waits before the failtimer expires.
Failtimer expiry action:	Displays the action taken when the failtimer expires: • Send-alert—Sends a critical message to the syslog when the failtimer expires. • Open-secondary-port—Opens the secondary port when the failtimer expires. Displays only for master nodes.
Preforwarding Timer interval:[1]	The configured value of the timer. This value is set internally by the EAPS software.
Last update:[1]	Displayed only for transit nodes; indicates the last time the transit node received a hello packet from the master node (identified by its MAC address).
EAPS Domain has … Controller Vlans:	Lists the assigned name and ID of the control VLAN.
EAPS Domain has … Protected Vlans:[2]	Lists the assigned names and VLAN IDs of all the protected VLANs configured on this EAPS domain.
Number of Protected Vlans:	The count of protected VLANs configured on this EAPS domain.

1. These fields apply only to transit nodes; they are not displayed for a master node.
2. This list is displayed when you use the `detail` keyword in the `show eaps` command.

11 Spanning Tree Protocol (STP)

This chapter covers the following topics:

Using the Spanning Tree Protocol (STP) functionality of the switch makes your network more fault tolerant. The following sections explain more about STP and the STP features supported by ExtremeWare.

NOTE

STP is a part of the 802.1d bridge specification defined by the IEEE Computer Society. To explain STP in terms used by the 802.1d specification, the switch will be referred to as a bridge.

Overview of the Spanning Tree Protocol

STP is a bridge-based mechanism for providing fault tolerance on networks. STP allows you to implement parallel paths for network traffic, and ensure that:

- Redundant paths are disabled when the main paths are operational.
- Redundant paths are enabled if the main path fails.

NOTE

STP is not supported in conjunction with ESRP.

Spanning Tree Domains

The switch can be partitioned into multiple virtual bridges. Each virtual bridge can run an independent Spanning Tree instance. Each Spanning Tree instance is called a *Spanning Tree Domain* (STPD). Each STPD has its own root bridge and active path. After an STPD is created, one or more VLANs can be assigned to it.

NOTE

A VLAN can span multiple STPDs. However, on the Summit 400, there is a hardware limitation that restricts each physical port to a single STPD. If the Summit 400 port is already a member of an STPD, then that port cannot be in another VLAN that is in a different STPD, or not in a STPD at all.

The key points to remember when configuring VLANs and STP are:

- Each VLAN forms an independent broadcast domain.
- STP blocks paths to create a loop-free environment.
- When STP blocks a path, no data can be transmitted or received on the blocked port.
- Within any given STPD, all VLANs belonging to it use the same spanning tree.

If you delete a STPD, the VLANs that were members of that STPD are also deleted. You must remove all VLANs associated with the STP before deleting the STPD to preserve the VLAN configuration.

STPD Modes

An STPD has two modes of operation

- 802.1d mode

  Use this mode for backward compatibility with previous STP versions and for compatibility with third-party switches using IEEE standard 802.1d. When configured in this mode, all rapid configuration mechanisms are disabled.

- 802.1w mode

  Use this mode for compatibility with Rapid Spanning Tree (RSTP). When configured in this mode, all rapid configuration mechanisms are enabled. This mode is available for point-to-point links only.

  RSTP is enabled or disabled on a per STPD basis only. You do not enable RSTP on a per port basis.

  For more information about RSTP and RSTP features, see “Rapid Spanning Tree Protocol” on page 198.

By default, the:

- STPD operates in 802.1d mode
- Default device configuration contains a single STPD called *s0*
- Default VLAN is a member of STPD s0

To configure the mode of operation of an STPD, use the following command:

```
configure stpd <spanning tree name> mode [dot1d | dot1w]
```

All STP parameters default to the IEEE 802.1d values, as appropriate.

Port Modes

An STP port has two modes of operation:

- 802.1d mode

  This mode is used for backward compatibility with previous STP versions and for compatibility with third-party switches using IEEE standard 802.1d. BPDUs are sent untagged in 1D mode. Because of this, on any given physical interface there can be only *one* STPD running in 1D mode.

- PVST+ mode

  This mode implements PVST+ in compatibility with third-party switches running this version of STP. The STPDs running in this mode have a one-to-one relationship with VLANs, and send and process packets in PVST+ format.

These port modes are for STP ports, not for physical ports. The Summit 400 restricts each physical port to a single STPD.

STPD Identifier

An StpdID is used to identify each STP domain. You assign the StpdID when configuring the domain, and that VLAN cannot belong to another STPD.

An StpdID must be identical to the VLANid of one of the member VLANs in that STP domain.

NOTE

If an STPD contains at least one port not in 1D mode, the STPD must be configured with an StpdID.

STPD BPDU Tunneling

You can configure ExtremeWare to allow a BDPU to traverse a VLAN without being processed by STP, even if STP is enabled on the port. This is known as BPDU *tunneling*.

To enable and disable BPDU tunneling on a VLAN, use one of the following commands:

```
enable ignore-bpdu vlan <vlan name>
disable ignore-bpdu vlan <vlan name>
```

If you have a known topology and have switches outside of your network within your STPD, use this feature to keep the root bridge within your network.

Rapid Root Failover

ExtremeWare supports rapid root failover for faster STP failover recovery times in STP 802.1d mode. If the active root port link goes down ExtremeWare recalculates STP and elects a new root port. Rapid root failover allows the new root port to immediately begin forwarding, skipping the standard listening and learning phases. Rapid root failover occurs only when the link goes down, and not when there is any other root port failure, such as missing BPDUs.

The default setting is disabled. To enable rapid root failover, use the following command:

```
enable stpd <spanning tree name> rapid-root-failover
```

To display the configuration, use the following command:

```
show stpd {<spanning tree name> | detail}
```

STP Configurations

When you assign VLANs to an STPD, pay careful attention to the STP configuration and its effect on the forwarding of VLAN traffic.

This section describes two types of STP configurations:

- Basic STP
- A VLAN that spans multiple STPDs

Basic STP Configuration

This section describes a basic, 802.1D STP configuration. Figure 29 illustrates a network that uses VLAN tagging for trunk connections. The following four VLANs have been defined:

- *Sales* is defined on switch A, switch B, and switch M.
- *Personnel* is defined on switch A, switch B, and switch M.
- *Manufacturing* is defined on switch Y, switch Z, and switch M.
- *Engineering* is defined on switch Y, switch Z, and switch M.
- *Marketing* is defined on all switches (switch A, switch B, switch Y, switch Z, and switch M).

Two STPDs are defined:

- STPD1 contains VLANs *Sales* and *Personnel.*
- STPD2 contains VLANs *Manufacturing* and *Engineering.*

The VLAN *Marketing* is a member of both STPD1 and STPD2.

Figure 29: Multiple Spanning Tree Domains

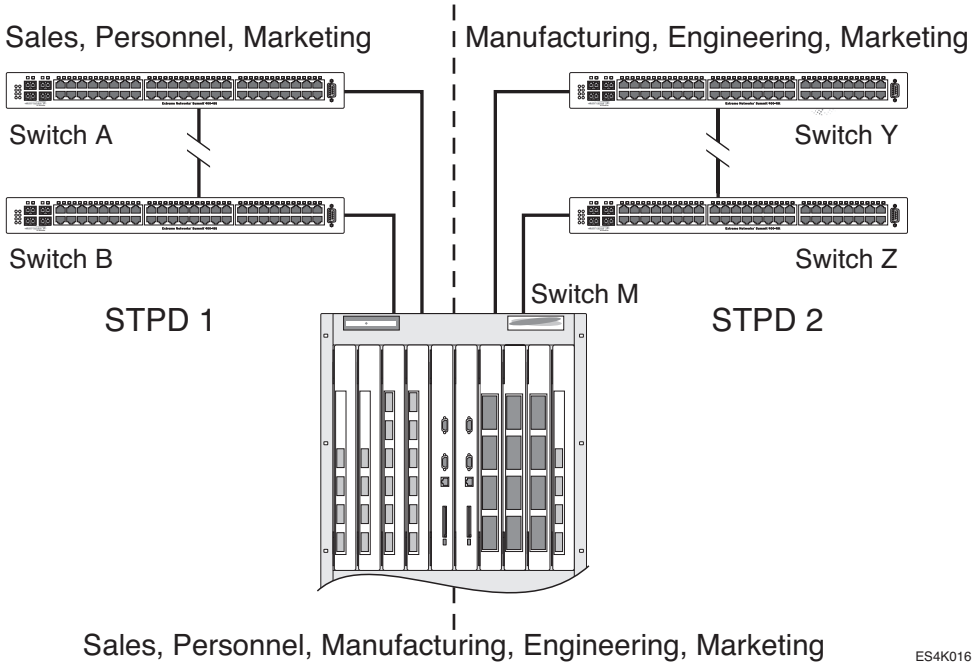

When the switches in this configuration start up, STP configures each STPD such that there are no active loops in the topology. STP could configure the topology in a number of ways to make it loop-free.

In Figure 29, the connection between switch 1 and switch 2 is put into blocking state, and the connection between switch Y and switch Z is put into blocking state. After STP converges, all the VLANs can communicate, and all bridging loops are prevented.

The VLAN *Marketing*, which has been assigned to both STPD1 and STPD2, communicates using all five switches. The topology has no loops, because STP has already blocked the port connection between switch A and switch B, and between switch Y and switch Z.

Within a single STPD, you must be extra careful when configuring your VLANs. Figure 30 illustrates a network that has been incorrectly set up using a single STPD so that the STP configuration disables the ability of the switches to forward VLAN traffic.

Figure 30: Tag-based STP configuration

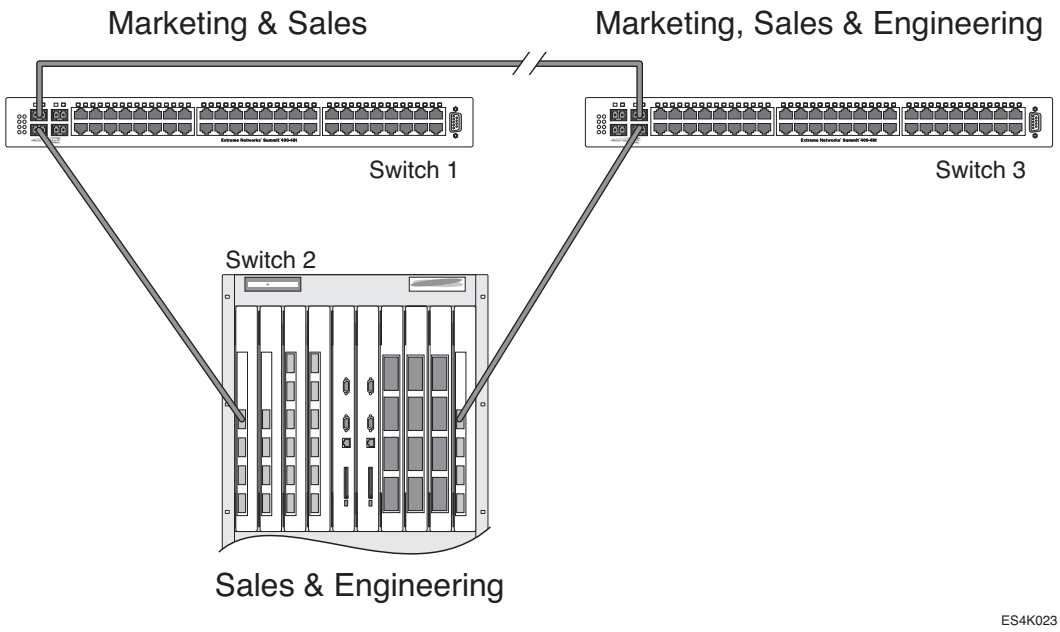

The tag-based network in Figure 30 has the following configuration:

- Switch 1 contains VLAN *Marketing* and VLAN *Sales.*
- Switch 2 contains VLAN *Engineering* and VLAN *Sales.*
- Switch 3 contains VLAN *Marketing*, VLAN *Engineering*, and VLAN *Sales.*
- The tagged trunk connections for three switches form a triangular loop that is not permitted in an STP topology.
- All VLANs in each switch are members of the same STPD.

STP can block traffic between switch 1 and switch 3 by disabling the trunk ports for that connection on each switch.

Switch 2 has no ports assigned to VLAN marketing. Therefore, if the trunk for VLAN marketing on switches 1 and 3 is blocked, the traffic for VLAN marketing will not be able to traverse the switches.

NOTE

If an STPD contains multiple VLANs, all VLANs must be configured on all ports in that domain, except for ports that connect to hosts (edge ports).

VLAN Spanning Multiple STPDs

Traditionally, the mapping from VLANs to STP instances have been one-to-one, or many-to-one. In both cases, a VLAN is wholly contained in a single instance. In practical deployment there are cases in which a one-to-many mapping is desirable. In a typical large enterprise network, for example, VLANs span multiple sites and/or buildings. Each site represents a redundant looped area. However, between any two sites the topology is usually very simple.

Alternatively, the same VLAN may span multiple large geographical areas (because they belong to the same enterprise) and may traverse a great many nodes. In this case, it is desirable to have multiple STP domains operating in a single VLAN, one for each looped area. The justifications include the following:

- The complexity of the STP algorithm increases, and performance drops, with the size and complexity of the network. The 802.1d standard specifies a maximum network diameter of 7 hops. By segregating a big VLAN into multiple STPDs, you reduce complexity and enhance performance.
- Local to each site, there may be other smaller VLANs that share the same redundant looped area with the large VLAN. Some STPDs must be created to protect those VLAN. The ability to partition VLANs allows the large VLAN to be “piggybacked” in those STPDs in a site-specific fashion.

Figure 31 has five domains. VLANs green, blue, brown, and yellow are local to each domain. VLAN red spans all of the four domains. Using a VLAN that spans multiple STPDS, you do not have to create a separate domain for VLAN red. Instead, VLAN red is “piggybacked” onto those domains local to other VLANs.

Figure 31: VLAN Spanning Multiple STPDs

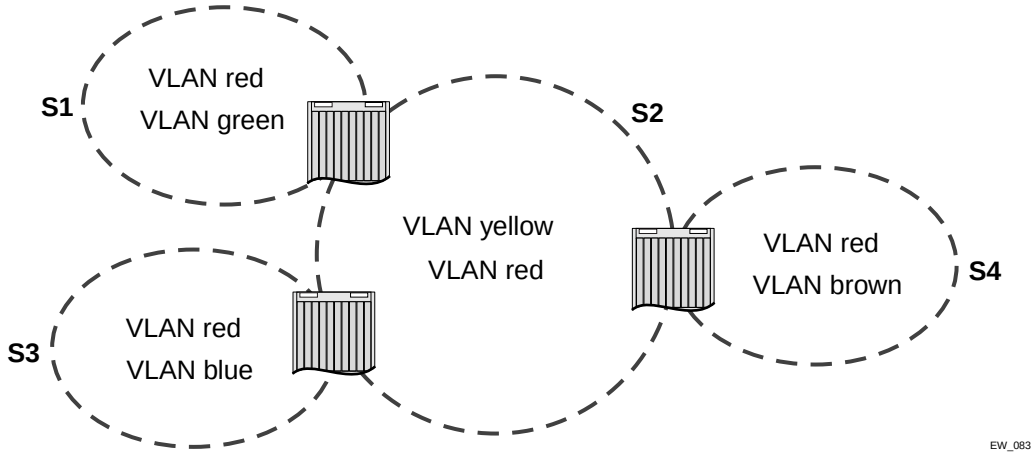

In addition, the configuration in Figure 31 has these features:

- Each site can be administered by a different organization or department within the enterprise. Having a site-specific STP implementation makes the administration more flexible and convenient.
- Between the sites the connection usually traverse distribution switches in ways that are known beforehand to be “safe” with STP. In other words, the looped areas are already well-defined.

Per-VLAN Spanning Tree

Switching products that implement Per-VLAN Spanning Tree (PVST) have been in existence for many years and are widely deployed. To support STP configurations that use PVST, ExtremeWare has an operational mode called PVST+.

NOTE

In this document, PVST and PVST+ are used interchangeably. PVST+ is an enhanced version of PVST that is interoperable with 802.1Q STP. The following discussions are in regard to PVST+, if not specifically mentioned.

STPD VLAN Mapping

Each VLAN participating in PVST+ must be in a separate STPD and the VLAN number must be the same as the STPD identifier (StpdID).As a result, PVST+ VLANs can not be partitioned.

This fact does not exclude other non-PVST+ VLANs from being grouped into the same STPD. A PVST+ VLAN can be joined by multiple non-PVST+ VLANs to be in the same STP domain.

Native VLAN

In PVST+, the native VLAN must be peered with default VLAN on Extreme devices, as both are the only VLAN allowed to send and receive untagged packets on the physical port.

Third-party PVST+ devices send VLAN 1 packets in a special manner. ExtremeWare does not support PVST+ for VLAN 1. Therefore, when the switch receives a packet for VLAN 1, the packet is dropped.

When a PVST+ instance is disabled, the fact that PVST+ uses a different packet format raises an issue. If the STPD also contains ports not in PVST+ mode, the flooded packet has an incompatible format with those ports. The packet is not recognized by the devices connected to those ports. Therefore, ExtremeWare has the following limitation:

- If an STPD contains both PVST+ and non-PVST+ ports, the STPD must not be disabled. Otherwise, the BPDUs are flooded in the format of the incoming STP port.

Rapid Spanning Tree Protocol

The Rapid Spanning Tree Protocol (RSTP; 802.1w) provides an enhanced spanning tree algorithm that improves the convergence speed of bridged networks. RSTP takes advantage of point-to-point links in the network and actively confirms that a port can safely transition to the forwarding state without relying on any timer configurations. If a network topology change or failure occurs, RSTP rapidly recovers network connectivity by confirming the change locally before propagating that change to other devices across the network. For broadcast links, there is no difference in convergence time between STP and RSTP.

RSTP supersedes legacy STP protocols, supports the existing STP parameters and configurations, and allows for seamless interoperability with legacy STP.

NOTE

RSTP is not supported in conjunction with ESRP.

RSTP Terms

Table 33 describes the terms associated with RSTP.

Table 33: RSTP Terms

Term	Description
root port	Provides the shortest path to the root bridge. All bridges except the root bridge, contain one root port. For more information about the root port, see “Port Roles” on page 199.
designated port	Provides the shortest path connection to the root bridge for the attached LAN segment. There is only one designated port on each LAN segment. For more information about the designated port, see “Port Roles” on page 199.
alternate port	Supplies an alternate path to the root bridge and the root port. For more information about the alternate port, see “Port Roles” on page 199.
backup port	Supports the designated port on the same attached LAN segment. Backup ports only exist when the bridge is connected as a self-loop or to a shared-media segment. For more information about the backup port, see “Port Roles” on page 199.
edge ports	Ports that connect to non-STP devices such as routers, endstations, and other hosts. Edge ports are not part of the RSTP configuration.
root bridge	The bridge with the best bridge identifier selected to be the root bridge. There is only one root bridge in the network. The root bridge is the only bridge in the network that does not have a root port.

RSTP Concepts

This section describes important RSTP concepts.

Port Roles

RSTP uses information from BPDUs to assign port roles for each LAN segment. Port roles are not user-configurable. Port role assignments are determined based on the following criteria:

- A unique bridge identifier (MAC address) associated with each bridge
- The path cost associated with each bridge port
- A port identifier associated with each bridge port

RSTP assigns one of four port roles to bridge ports in the network, as described in Table 34.

Table 34: RSTP port roles

Port Role	Description
Root	Provides the shortest path to the root bridge. There is only one root port per bridge; the root bridge does not have a root port. If a bridge has two or more ports with the same path cost, the port with the best port identifier becomes the root port.
Designated	Provides the shortest path connection to the root bridge for the attached LAN segment. To prevent loops in the network, there is only one designated port on each LAN segment. To select the designated port, all bridges that are connected to a particular segment listen to each other's BPDUs and agree on the bridge sending the best BPDU. The corresponding port on that bridge becomes the designated port. If there are two or more ports connected to the LAN, the port with the best port identifier (lowest MAC address) becomes the designated port.
Alternate	Provides an alternate path to the root bridge and the root port.
Backup	Supports the designated port on the same attached LAN segment. Backup ports only exist when the bridge is connected as a self-loop or to a shared-media segment.

When RSTP stabilizes, all:

- Root ports and designated ports are in the forwarding state
- Alternate ports and backup ports are in the blocking state

RSTP makes the distinction between the alternate and backup port roles to describe the rapid transition of the alternate port to the forwarding state if the root port fails.

Ports that connect to non-STP devices are edge ports. Edge ports do not participate in RSTP, and their role is not confirmed. Edge ports immediately enter the forwarding state.

Link Types

You can configure the link type of a port in an STPD. RSTP tries to rapidly move designated point-to-point links into the forwarding state when a network topology change or failure occurs. For rapid convergence to occur, the port must be configured as a point-to-point link.

Table 35 describes the link types.

Table 35: RSTP link types

Port Role	Description
Auto	Specifies the switch to automatically determine the port link type. An auto link behaves like a point-to-point link if the link is in full duplex mode or if link aggregation is enabled on the port. Otherwise, the link behaves like a broadcast link used for 802.1w configurations.
Edge	Specifies a port that does not have a bridge attached. An edge port is placed and held in the STP forwarding state unless a BPDU is received by the port.
Broadcast	Specifies a port attached to a LAN segment with more than two bridges. A port with a broadcast link type cannot participate in rapid reconfiguration. By default, all ports are broadcast links.
Point-to-point	Specifies a port attached to a LAN segment with only two bridges. A port with port-to-port link type can participate in rapid reconfiguration. Used for 802.1w configurations.

Configuring Link Types. By default, all ports are broadcast links. To configure the ports in an STPD, use the following command:

```
configure stpd <spanning tree name> ports link-type [auto | edge | broadcast |
point-to-point] <portlist>
```

- `auto`—Configures the ports as auto links. If the link is in full duplex mode, or if link aggregation is enabled on the port, an auto link behaves like a point-to-point link.
- `edge`—Configures the ports as edge ports.
- `point-to-point`—Configures the ports for an RSTP environment.

To display detailed information about the ports in an STPD, use the following command:

```
show stpd <spanning tree name> ports <portlist> {detail}
```

RSTP Timers

For RSTP to rapidly recover network connectivity, RSTP requires timer expiration. RSTP derives many of the timer values from the existing configured STP timers to meet its rapid recovery requirements rather than relying on additional timer configurations. Table 36 describes the user configurable timers, and Table 37 describes the timers that are derived from other timers and not user configurable.

Table 36: User configurable timers

Timer	Description
Hello	The root bridge uses the hello timer to send out configuration BPDUs through all of its forwarding ports at a pre-determined, regular time interval. The default value is 2 seconds. The range is 1 to 10 seconds.
Forward delay	A port moving from the blocking state to the forwarding state uses the forward delay timer to transition through the listening and learning states. In RSTP, this timer complements the rapid configuration behavior. If none of the rapid rules are in effect, the port uses legacy STP rules to move to the forwarding state. The default is 15 seconds. The range is 4 to 30 seconds.

Table 37: Derived timers

Timer	Description
TCN	The root port uses the TCN timer when it detects a change in the network topology. The TCN timer stops when the topology change timer expires or upon receipt of a topology change acknowledgement. The default value is the same as the value for the bridge hello timer.
Topology Change	The topology change timer determines the total time it takes the forwarding ports to send configuration BPDUs. The default value for the topology change timer depends upon the mode of the port. • 1d mode—The sum of the forward delay timer (default value is 15 seconds; range of 4 to 30 seconds) and the max age timer (default value is 20 seconds; range of 6 to 40 seconds). • 1w mode—Double the hello timer (default value is 4 seconds)
Message age	A port uses the message age timer to time out receiving BPDUs. When a port receives a superior or equal BPDU, the timer restarts. When the timer expires, the port becomes a designated port and a configuration update occurs. If the bridge operates in 1w mode and receives an inferior BPDU, the timer expires early. The default value is the same as the STPD bridge max age parameter.

Table 37: Derived timers (Continued)

Timer	Description
Hold	A port uses the hold timer to restrict the rate that successive BPDUs can be sent. The default value is the same as the value for the bridge hello timer.
Recent backup	The timer starts when a port leaves the backup role. When this timer is running, the port cannot become a root port. The default value is double the hello time (4 seconds).
Recent root	The timer starts when a port leaves the root port role. When this timer is running, another port cannot become a root port unless the associated port is put into the blocking state. The default value is the same as the forward delay time.

The Protocol migration timer is neither user-configurable nor derived; it has a set value of 3 seconds. The timer starts when a port transitions from STP (802.1d) mode to RSTP (802.1w) mode and vice versa. This timer must expire before further mode transitions can occur.

RSTP Operation

In an RSTP environment, there are two bridges on a point-to-point link LAN segment. A switch that considers itself the unique, designated bridge for the attached LAN segment sends a “propose” message to the other bridge to request a confirmation of its role. The other bridge on that LAN segment replies with an “agree” message if they agree with the proposal. The receiving bridge immediately moves its designated port into the forwarding state.

Before a bridge replies with an “agree” message, it reverts all of its designated ports into the blocking state. This introduces a temporary partition into the network. The bridge then sends another “propose” message on all of its designated ports for further confirmation. Since all of the connections are blocked, the bridge immediately sends an “agree” message to unblock the proposing port without having to wait for further confirmations to come back or without the worry of temporary loops.

Beginning with the root bridge, each bridge in the network engages in the exchange of “propose” and “agree” messages until they reach the edge ports. Edge ports connect to non-STP devices and do not participate in RSTP. Their role does not need to be confirmed. If an edge port receives a BPDU, it enters an inconsistency state. An inconsistency state puts the edge port into the blocking state and starts the message age timer. Every time the edge port receives a BPDU, the message age timer restarts. The edge port remains in the blocking state until no further BPDUs are received and the message age timer expires.

RSTP attempts to transition root ports and designated ports to the forwarding state and alternate ports and backup ports to the blocking state as rapidly as possible.

A port transitions to the forwarding state if any of the following is true. The port:

- Has been in either a root or designated port role long enough that the spanning tree information supporting this role assignment has reached all of the bridges in the network.

  NOTE

  RSTP is backward compatible with STP, so if a port does not move to the forwarding state with any of the RSTP rapid transition rules, a forward delay timer starts and STP behavior takes over.

- Is now a root port and no other ports have a recent role assignment that contradicts with its root port role.

- Is a designated port and attaches to another bridge by a point-to-point link and receives an “agree” message from the other bridge port.
- Is an edge port.

  An edge port is a port connected to a non-STP device and is in the forwarding state.

The preceding sections provide more information about RSTP behavior.

Root Port Rapid Behavior

In Figure 32, the diagram on the left displays the initial network topology with a single bridge having the following:

- Two ports connected to a shared LAN segment
- One port is the designated port
- One port is the backup port

The diagram on the right displays a new bridge that:

- Is connected to the LAN segment
- Has a superior STP bridge priority
- Becomes the root bridge and sends a BPDU to the LAN that is received by both ports on the old bridge

Figure 32: Example of root port rapid behavior

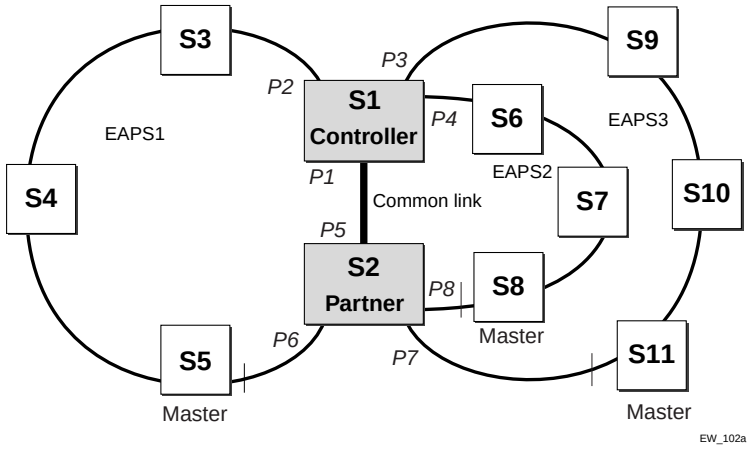

If the backup port receives the BPDU first, STP processes this packet and temporarily elects this port as the new root port while the designated port’s role remains unchanged. If the new root port is immediately put into the forwarding state, there is a loop between these two ports.

To prevent this type of loop from occurring, the recent backup timer starts. The root port transition rule does not allow a new root port to be in the forwarding state until the recent backup timer expires.

Another situation may arise if you have more than one bridge, and you lower the port cost for the alternate port which makes it the new root port. The previous root port is now an alternate port. Depending on your STP implementation, STP may set the new root port to the forwarding state before setting the old root port to the blocking state. This may cause a loop.

To prevent this type of loop from occurring, the recent root timer starts when the port leaves the root port role. The timer stops if the port enters the blocking state. RSTP requires that the recent root timer stops on the previous root port before the new root port can enter the forwarding state.

Designated Port Rapid Behavior

When a port becomes a new designated port, or the STP priority changes on an existing designated port, the port becomes an *unsynced* designated port. In order for an unsynced designated port to rapidly move into the forwarding state, the port must propose a confirmation of its role on the attached LAN segment, unless the port is an edge port. Upon receiving an “agree” message, the port immediately enters the forwarding state.

If the receiving bridge does not agree and it has a superior STP priority, the receiving bridge replies with its own BPDU. Otherwise, the receiving bridge keeps silent and the proposing port enters the forwarding state and starts the forward delay timer.

The link between the new designated port and the LAN segment must be a point-to-point link. If there is a multi-access link, the “propose” message is sent to multiple recipients. If only one of the recipients agrees with the proposal, it is possible for the port to erroneously enter the forwarding state after receiving a single “agree” message.

Receiving Bridge Behavior

The receiving bridge must decide whether or not to accept a proposal from a port. Upon receiving a proposal for a root port, the receiving bridge:

- Processes the BPDU and computes the new STP topology
- Synchronizes all of the designated ports if the receiving port is the root port of the new topology
- Puts all unsynced, designated ports into the blocking state
- Sends down further “propose” messages
- Sends back an “agree” message through the root port

If the receiving bridge receives a proposal for a designated port, the bridge replies with its own BPDU. If the proposal is for an alternate or backup port, the bridge keeps silent.

Propagating Topology Change Information

When a change occurs in the topology of the network, such events are communicated through the network.

In an RSTP environment, only non-edge ports entering the forwarding state cause a topology change. A loss of network connectivity is not considered a topology change; however, a gain in network connectivity needs to be communicated. When an RSTP bridge detects a topology change, it starts the topology change timer, sets the topology change flag on its BPDUs, floods all of the forwarding ports in the network (including the root ports), and flushes the learned MAC address entries.

Rapid Reconvergence

This section describes the RSTP rapid behavior following a topology change. In this example, the bridge priorities are assigned based on the order of their alphabetical letters; bridge A has a higher priority than bridge F.

Suppose we have a network, as shown in Figure 33, with six bridges (bridge A through bridge F) where the following is true:

- Bridge A is the root bridge
- Bridge D contains an alternate port in the blocking state
- All other ports in the network are in the forwarding state

Figure 33: Initial network configuration

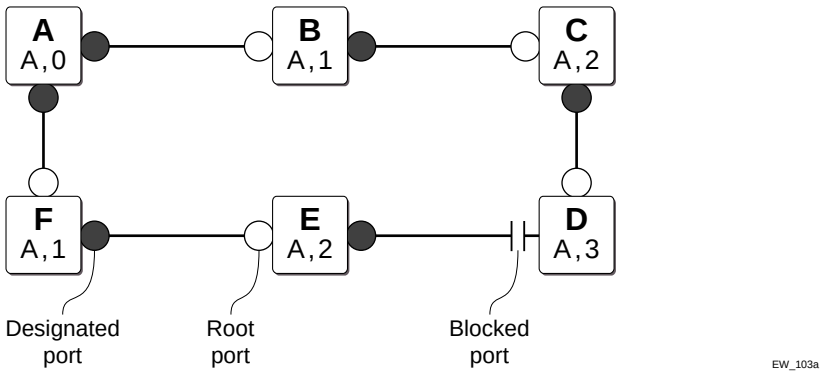

The preceding steps describe how the network reconverges.

1 If the link between bridge A and bridge F goes down, bridge F detects the root port is down. At this point, bridge F:

- Immediately deletes that port from the STP
- Performs a configuration update

After the configuration update, bridge F:

- Considers itself the new root bridge
- Sends a BPDU message on its designated port to bridge E

Figure 34: Down link detected

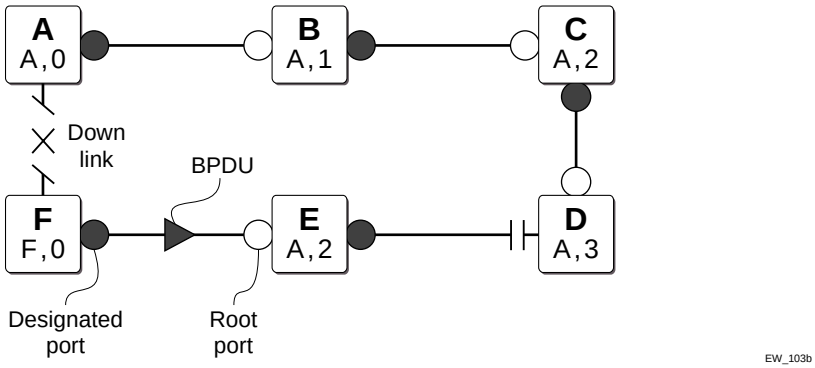

2 Bridge E believes that bridge A is the root bridge. When bridge E receives the BPDU on its root port from bridge F, bridge E:

- Determines that it received an inferior BPDU.
- Immediately begins the max age timer on its root port
- Performs a configuration update

After the configuration update, bridge E:

- Regards itself as the new root bridge
- Sends BPDU messages on both of its root ports to bridges F and D, respectively

Figure 35: New root bridge selected

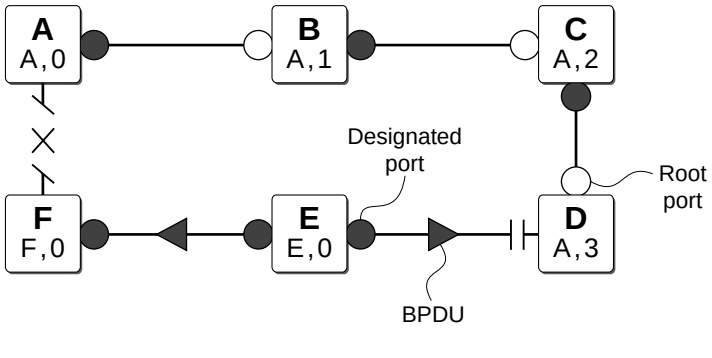

3 When bridge F receives the superior BPDU and configuration update from bridge E, bridge F:

- Decides that the receiving port is the root port
- Determines that bridge E is the root bridge.

Figure 36: Communicating new root bridge status to neighbors

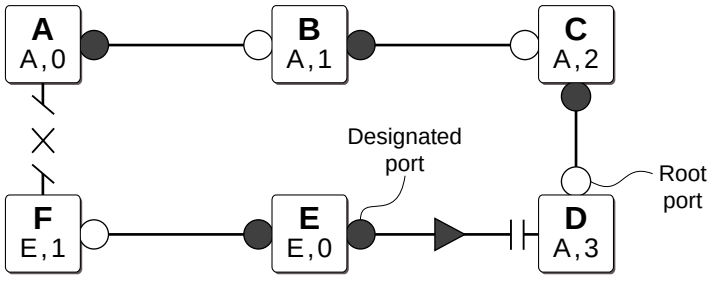

4 Bridge D believes that bridge A is the root bridge. When bridge D receives the BPDU from bridge E on its alternate port, bridge D:

- Immediately begins the max age timer on its alternate port
- Performs a configuration update

After the configuration update, bridge D:

- Moves the alternate port to a designated port
- Sends a “propose” message to bridge E to solicit confirmation of its designated role and to rapidly move the port into the designated state

Figure 37: Sending a propose message to confirm a port role

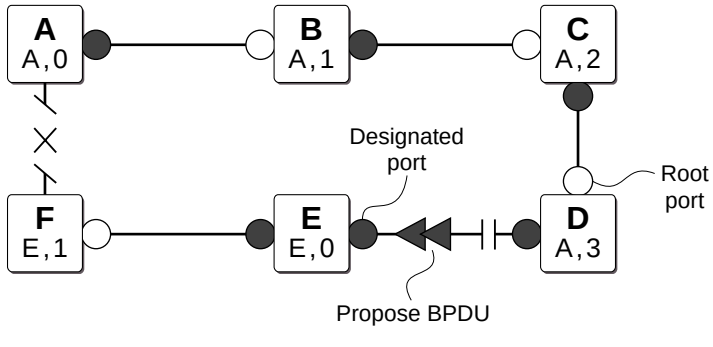

5 Upon receiving the proposal, bridge E:

- Performs a configuration update
- Changes its receiving port to a root port

  The existing designated port enters the blocking state

Bridge E then sends:

- A “propose” message to bridge F
- An “agree” message from its root port to bridge D.

Figure 38: Communicating port status to neighbors

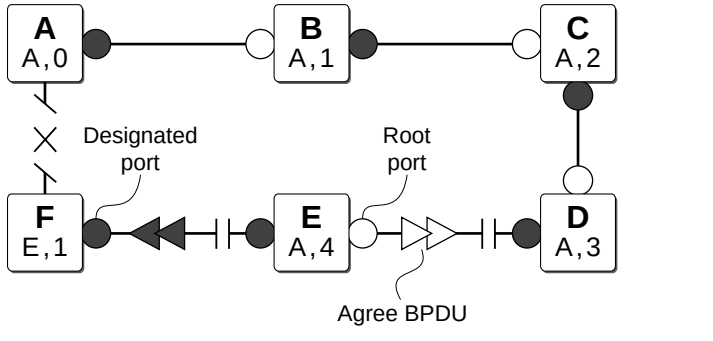

6 To complete the topology change, the following occurs:

- Bridge D moves the port that received the agree message into the forwarding state
- Bridge F confirms that its receiving port (the port that received the “propose” message) is the root port, and immediately replies with an “agree” message to bridge E to unblock the proposing port

Figure 39: Completing the topology change

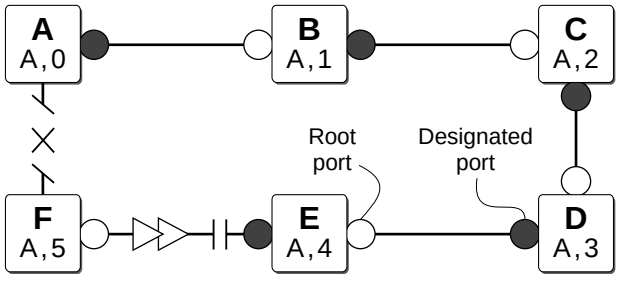

Figure 40 displays the new topology.

Figure 40: Final network configuration

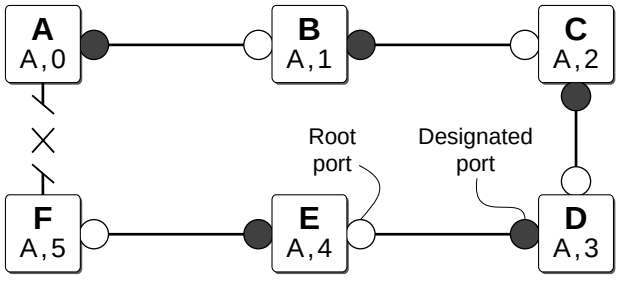

Compatibility With STP (802.1d)

RSTP interoperates with legacy STP protocols; however, the rapid convergence benefits are lost when interacting with legacy STP bridges.

Each RSTP bridge contains a port protocol migration state machine to ensure that the ports in the STPD operate in the correct, configured mode. The state machine is a protocol entity within each bridge configured to run in 802.1w mode. For example, a compatibility issue occurs if you configure 802.1w mode and the bridge receives an 802.1d BPDU on a port. The receiving port starts the protocol migration timer and remains in 802.1d mode until the bridge stops receiving 802.1d BPDUs. Each time the bridge receives an 802.1d BPDU, the timer restarts. When the port migration timer expires, no more 802.1d BPDUs have been received and the bridge returns to its configured setting, 802.1w mode.

STP Rules and Restrictions

This section summarizes the rules and restrictions for configuring STP.

- The StpdID must be the VLANid of one of its member VLANs, and that VLAN can not be partitioned.
- A default VLAN can not be partitioned. If a VLAN traverses multiple STP domains, the VLAN must be tagged.
- An STPD can carry, at most, one VLAN running in PVST+ mode, and its StpdID must be identical with that VLANid. In addition, the PVST+ VLAN can not be partitioned.
- The default VLAN of a PVST+ port must be identical with the native VLAN on the PVST+ device connected to that port.
- If a port supports 802.1w-STPD, then the port must be configured with a default VLAN. If not, the BPDUs for that STPD are not flooded when the STPD is disabled.
- If an STPD contains both PVST+ and non-PVST+ ports, it must be enabled. If it is disable, the BPDUs are flooded in the format of the incoming STP port, which may be incompatible with those of the connected devices.

Configuring STP on the Switch

To configure basic STP, follow these steps:

1 Create one or more STP domains using the following command:

```
create stpd <name>
```

NOTE

STPD, VLAN, and QoS profile names must all be unique. For example, a name used to identify a VLAN cannot be used when you create an STPD or a QoS profile.

2 Add one or more VLANs to the STPD using the following command:

```
configure stpd <spanning tree name> add vlan <vlan name> {ports <portlist> [dot1d
| pvst-plus]}
```

3 Enable STP for one or more STP domains using the following command:

```
enable stpd {<spanning tree name>}
```

After you have created the STPD, you can optionally configure STP parameters for the STPD.

NOTE

You should not configure any STP parameters unless you have considerable knowledge and experience with STP. The default STP parameters are adequate for most networks.

The following parameters can be configured on each STPD:

- Hello time
- Forward delay
- Max age

- Bridge priority
- StpdID

The following parameters can be configured on each port:

- Path cost
- Port priority
- Port mode

NOTE

The device supports the RFC 1493 Bridge MIB. Parameters of only the s0 default STPD are accessible through this MIB.

NOTE

If an STPD contains at least one port not in dot1D mode, the STPD must be configured with an StpdID.

STP Configuration Examples

This section provides three configuration examples:

- Basic 802.1d STP
- RSTP 802.1w

Basic 802.1d Configuration Example

The following example creates and enables an STPD named *Backbone_st*. It assigns the *Manufacturing* VLAN to the STPD. It disables STP on ports 1 through 7, and port 12.

```
create stpd backbone_st
configure stpd backbone_st add vlan manufacturing
enable stpd backbone_st
disable stpd backbone_st port 1-7,12
```

RSTP 802.1w Configuration Example

Figure 41 is an example of a network with multiple STPDs that can benefit from RSTP. For RSTP to work, you need to do the following:

- Create an STPD
- Configure the mode of operation for the STPD
- Create the VLANs and assign the ports
- Add the VLANs to the STPD
- Configure the port link types
- Enable STP

Figure 41: RSTP example

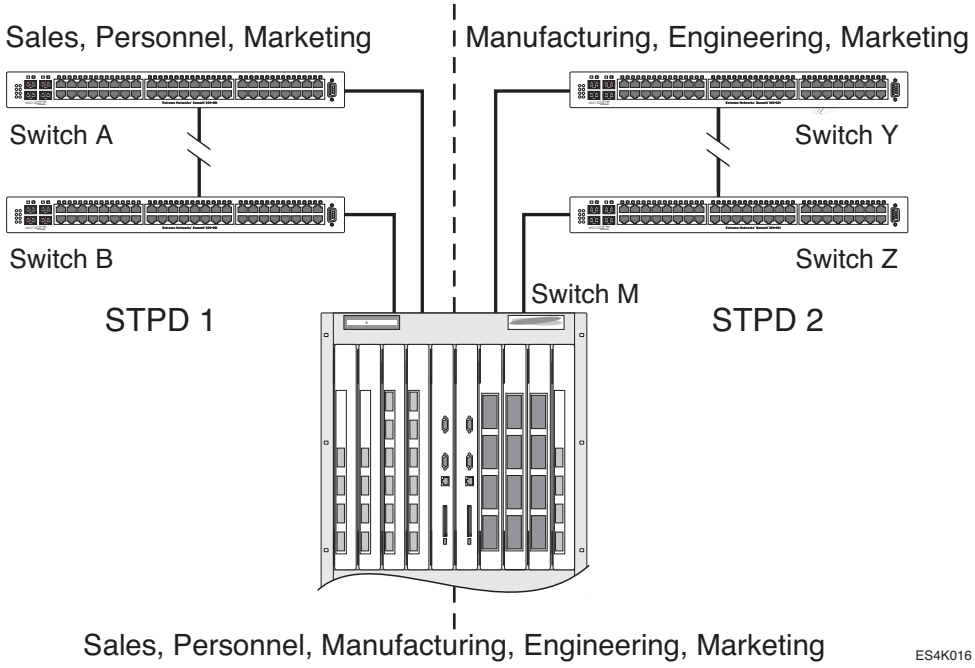

In this example, the commands configure switch A in STPD1 for rapid reconvergence. Use the same commands to configure each switch and STPD in the network.

```
create stpd stpd1
configure stpd stpd1 mode dot1w

create vlan sales
create vlan personnel
create vlan marketing
configure vlan sales add ports 1,2 tagged
configure vlan personnel add ports 1,2 tagged
configure vlan marketing add ports 1,2 tagged

configure stpd stpd1 add vlan sales
configure stpd stpd1 add vlan personnel
configure stpd stpd1 add vlan marketing

configure stpd stpd1 ports link-type point-to-point 1,2

enable stpd stpd1
```

Displaying STP Settings

To display STP settings, use the following command:

```
show stpd {<spanning tree name> | detail}
```

This command displays the following information:

- STPD name
- STPD state
- STPD mode of operation
- Rapid Root Failover
- Tag
- Ports
- Active VLANs
- Bridge Priority
- Bridge ID
- Designated root
- STPD configuration information

To display the STP state of a port, use the following command:

```
show stpd <spanning tree name> ports <portlist> {detail}
```

This command displays the following information:

- STPD port configuration
- STPD port mode of operation
- STPD path cost
- STPD priority
- STPD state (root bridge, and so on)
- Port role (root bridge, edge port, etc.)
- STPD port state (forwarding, blocking, and so on)
- Configured port link type
- Operational port link type

If you have a VLAN that spans multiple STPDs, use the `show vlan <vlan name> stpd` command to display the STP configuration of the ports assigned to that specific VLAN.

The command displays the following:

- STPD port configuration
- STPD port mode of operation
- STPD path cost
- STPD priority
- STPD state (root bridge, and so on)
- Port role (root bridge, edge port, etc.)

- STPD port state (forwarding, blocking, and so on)
- Configured port link type
- Operational port link type

12 IP Unicast Routing

This chapter describes the following topics:

This chapter assumes that you are already familiar with IP unicast routing. If not, refer to the following publications for additional information:

- RFC 1256—*ICMP Router Discovery Messages*
- RFC 1812—*Requirements for IP Version 4 Routers*

NOTE

For more information on interior gateway protocols, see Chapter 13.

Overview of IP Unicast Routing

The switch provides full layer 3, IP unicast routing. It exchanges routing information with other routers on the network using either the Routing Information Protocol (RIP) or the Open Shortest Path First (OSPF) protocol. The switch dynamically builds and maintains a routing table, and determines the best path for each of its routes.

Each host using the IP unicast routing functionality of the switch must have a unique IP address assigned. In addition, the default gateway assigned to the host must be the IP address of the router interface.

Router Interfaces

The routing software and hardware routes IP traffic between router interfaces. A router interface is simply a VLAN that has an IP address assigned to it.

As you create VLANs with IP addresses belonging to different IP subnets, you can also choose to route between the VLANs. Both the VLAN switching and IP routing function occur within the switch.

NOTE

Each IP address and mask assigned to a VLAN must represent a unique IP subnet. You cannot configure the IP address belonging to the same subnet on different VLANs.

In Figure 42, a switch is depicted with two VLANs defined; *Finance* and *Personnel*. Port 8-15 are assigned to *Finance*; ports 24-48 are assigned to *Personnel*. *Finance* belongs to the IP network 192.207.35.0; the router interface for *Finance* is assigned the IP address 192.206.35.1. *Personnel* belongs to the IP network 192.207.36.0; its router interface is assigned IP address 192.207.36.1. Traffic within each VLAN is switched using the Ethernet MAC addresses. Traffic between the two VLANs is routed using the IP addresses.

Figure 42: Routing between VLANs

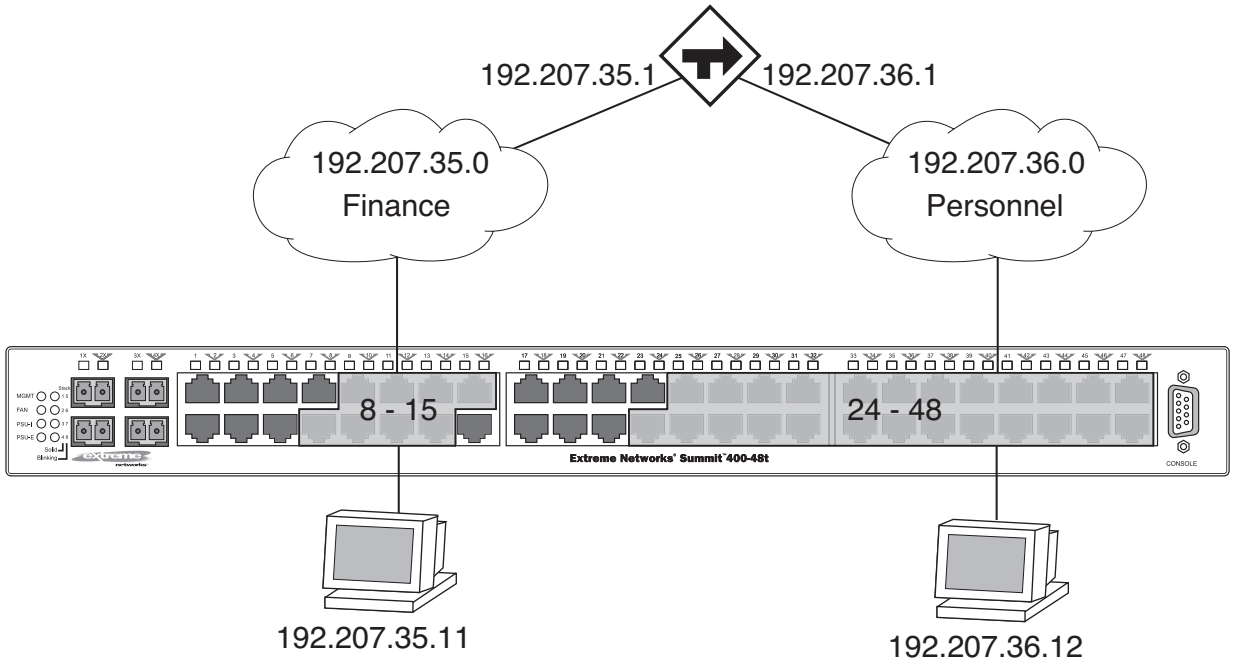

Populating the Routing Table

The switch maintains an IP routing table for both network routes and host routes. The table is populated from the following sources:

- Dynamically, by way of routing protocol packets or by ICMP redirects exchanged with other routers
- Statically, by way of routes entered by the administrator
  - Default routes, configured by the administrator
  - Locally, by way of interface addresses assigned to the system
  - By other static routes, as configured by the administrator

NOTE

If you define a default route, and subsequently delete the VLAN on the subnet associated with the default route, the invalid default route entry remains. You must manually delete the configured default route.

Dynamic Routes

Dynamic routes are typically learned by way of RIP or OSPF. Routers that use RIP or OSPF exchange information in their routing tables in the form of advertisements. Using dynamic routes, the routing table contains only networks that are reachable.

Dynamic routes are aged out of the table when an update for the network is not received for a period of time, as determined by the routing protocol.

Static Routes

Static routes are manually entered into the routing table. Static routes are used to reach networks not advertised by routers.

Static routes can also be used for security reasons, to control which routes you want advertised by the router. You can decide if you want all static routes to be advertised, using one of the following commands:

- `enable rip exportstatic` or `disable rip exportstatic`
- `enable ospf export static [cost <metric> [ase-type-1 | ase-type-2] {tag <number>}]` or `disable ospf export [direct | rip | static]`

The default setting is disabled. Static routes are never aged out of the routing table.

A static route must be associated with a valid IP subnet. An IP subnet is associated with a single VLAN by its IP address and subnet mask. If the VLAN is subsequently deleted, the static route entries using that subnet must be deleted manually.

Multiple Routes

When there are multiple, conflicting choices of a route to a particular destination, the router picks the route with the longest matching network mask. If these are still equal, the router picks the route using the following criteria (in the order specified):

- Directly attached network interfaces
- ICMP redirects

- Static routes
- Directly attached network interfaces that are not active.

NOTE

If you define multiple default routes, the route that has the lowest metric is used. If multiple default routes have the same lowest metric, the system picks one of the routes.

You can also configure *blackhole* routes—traffic to these destinations is silently dropped.

IP Route Sharing

IP route sharing allows multiple equal-cost routes to be used concurrently. IP route sharing can be used with static routes or with OSPF routes. In OSPF, this capability is referred to as *equal cost multipath* (ECMP) routing. To use IP route sharing, use the following command:

```
enable iproute sharing
```

Next, configure static routes and/or OSPF as you would normally. ExtremeWare supports unlimited route sharing across static routes and up to 12 ECMP routes for OSPF.

Route sharing is useful only in instances where you are constrained for bandwidth. This is typically not the case using Extreme switches. Using route sharing makes router troubleshooting more difficult because of the complexity in predicting the path over which the traffic will travel.

Subnet-Directed Broadcast Forwarding

You can enable or disable the hardware forwarding of subnet-directed broadcast IP packets. This allows the switch to forward subnet-directed broadcast packets at wire-speed.

To enable or disable hardware forwarding, use one the following commands:

```
[enable | disable] ipforwarding [vlan <vlan_name>]
```

The entries are added to the IP forwarding table as standard entries and you can view them using the `show ipfdb` command.

You can also configure the VLAN router interface to either forward and process all subnet-directed broadcast packets, or to simply forward these packets after they have been added to the IP forwarding database. The latter option allows you to improve CPU forwarding performance by having upper layers, such as UDP and TCP, ignore broadcast packet processing (for example, if the packets have IP-options configured).

To enable or disable broadcast packet processing, use the following command:

```
[enable | disable] ipforwarding ignore-broadcast vlan <vlan_name>
```

Using these commands together, you can achieve a 100% reduction on the Summit switches.

Proxy ARP

Proxy Address Resolution Protocol (ARP) was first invented so that ARP-capable devices could respond to ARP Request packets on behalf of ARP-incapable devices. Proxy ARP can also be used to achieve

router redundancy and simplify IP client configuration. The switch supports proxy ARP for this type of network configuration. The section describes some example of how to use proxy ARP with the switch.

ARP-Incapable Devices

To configure the switch to respond to ARP Requests on behalf of devices that are incapable of doing so, you must configure the IP address and MAC address of the ARP-incapable device using the use the following command:

```
configure iparp add proxy <ip address> {<mask>} {<mac_address>} {always}
```

Once configured, the system responds to ARP Requests on behalf of the device as long as the following conditions are satisfied:

- The valid IP ARP Request is received on a router interface.
- The target IP address matches the IP address configured in the proxy ARP table.
- The proxy ARP table entry indicates that the system should always answer this ARP Request, regardless of the ingress VLAN (the `always` parameter must be applied).

Once all the proxy ARP conditions are met, the switch formulates an ARP Response using the configured MAC address in the packet.

Proxy ARP Between Subnets

In some networks, it is desirable to configure the IP host with a wider subnet than the actual subnet mask of the segment. Proxy ARP can be used so that the router answers ARP Requests for devices outside of the subnet. As a result, the host communicates as if all devices are local. In reality, communication with devices outside of the subnet are proxied by the router.

For example, an IP host is configured with a class B address of 100.101.102.103 and a mask of 255.255.0.0. The switch is configured with the IP address 100.101.102.1 and a mask of 255.255.255.0. The switch is also configured with a proxy ARP entry of IP address 100.101.0.0 and mask 255.255.0.0, *without* the `always` parameter.

When the IP host tries to communicate with the host at address 100.101.45.67, the IP hosts communicates as if the two hosts are on the same subnet, and sends out an IP ARP Request. The switch answers on behalf of the device at address 100.101.45.67, using its own MAC address. All subsequent data packets from 100.101.102.103 are sent to the switch, and the switch routes the packets to 100.101.45.67.

Relative Route Priorities

Table 38 lists the relative priorities assigned to routes depending upon the learned source of the route.

NOTE

Although these priorities can be changed, do not attempt any manipulation unless you are expertly familiar with the possible consequences.

Table 38: Relative Route Priorities

Route Origin	Priority
Direct	10
BlackHole	50
Static	1100
ICMP	1200
OSPFIntra	2200
OSPFInter	2300
RIP	2400
OSPFExtern1	3200
OSPFExtern2	3300
BOOTP	5000

To change the relative route priority, use the following command:

```
configure iproute priority [rip | bootp | icmp | static | ospf-intra | ospf-inter |
ospf-as-external | ospf-extern1 | ospf-extern2] <priority>
```

Configuring IP Unicast Routing

This section describes the commands associated with configuring IP unicast routing on the switch. To configure routing, follow these steps:

1. Create and configure two or more VLANs.
2. Assign each VLAN that will be using routing an IP address using the following command:

   ```
   configure vlan <vlan name> ipaddress <ipaddress> {<netmask> | <mask length>}
   ```

   Ensure that each VLAN has a unique IP address.
3. Configure a default route using the following command:

   ```
   configure iproute add default <gateway> {<metric>}
   ```

   Default routes are used when the router has no other dynamic or static route to the requested destination.
4. Turn on IP routing for one or all VLANs using the following command:

   ```
   enable ipforwarding {[broadcast | ignore-broadcast]}{vlan <vlan name>}
   ```
5. Turn on RIP or OSPF using one of the following commands:

   ```
   enable ripp
   ```

   ```
   enable ospff
   ```

For more information on configuring RIPP and OSPF, see “Interior Gateway Protocols” on page 227.

Verifying the IP Unicast Routing Configuration

Use the `show iproute` command to display the current configuration of IP unicast routing for the switch, and for each VLAN. The `show iproute` command displays the currently configured routes, and includes how each route was learned.

Additional verification commands include:

- `show iparp`—Displays the IP ARP table of the system.
- `show ipfdb`—Displays the hosts that have been transmitting or receiving packets, and the port and VLAN for each host.
- `show ipconfig`—Displays configuration information for one or more VLANs.

Routing Configuration Example

Figure 43 illustrates a switch that has three VLANs defined as follows:

- *Finance*
  - Contain ports 5 and 6.
  - IP address 192.207.35.1.
- *Personnel*
  - Contain ports 21 and 22.
  - IP address 192.207.36.1.

Figure 43: Unicast routing configuration example

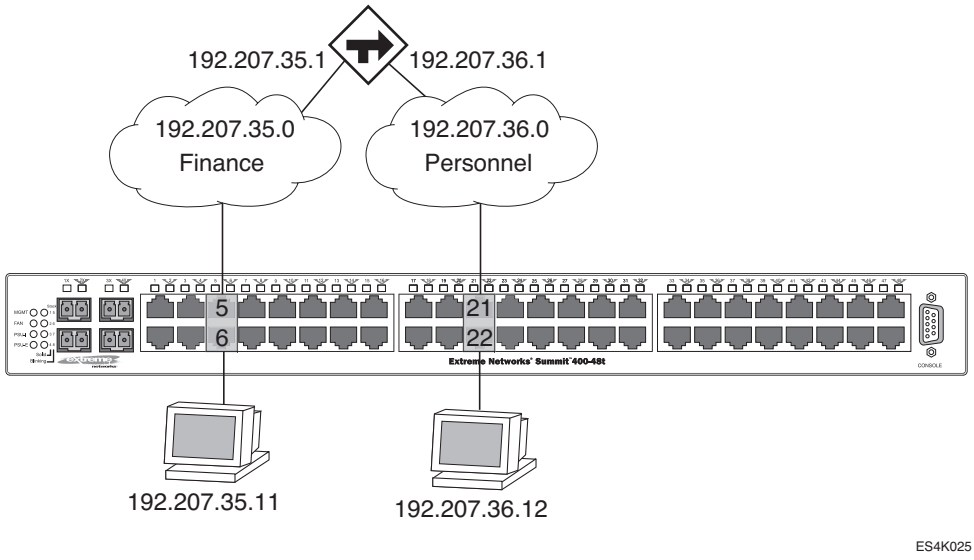

In this configuration, all IP traffic from stations connected to ports 5 and 6 have access to the switch by way of the VLAN *Finance*. Ports 21 and 22 reach the switch by way of the VLAN *Personnel*..

The example in Figure 43 is configured as follows:

```
create vlan Finance
```

```
create vlan Personnel

config Finance add port 5,6
config Personnel add port 21,22

config Finance ipaddress 192.207.35.1
config Personnel ipaddress 192.207.36.1

config rip add vlan Finance
config rip add vlan Personnel

enable ipforwarding
enable rip
```

ICMP Packet Processing

As ICMP packets are routed or generated, you can take various actions to control distribution. For ICMP packets typically generated or observed as part of the routing function, you can assert control on a per-type, per-VLAN basis. You would alter the default settings for security reasons: to restrict the success of tools that can be used to find an important application, host, or topology information. The controls include the disabling of transmitting ICMP messages associated with unreachables, port-unreachables, time-exceeded, parameter-problems, redirects, time-stamp, and address-mask requests.

To enable or disable the generation of an ICMP address-mask reply on one or all VLANs, use the following commands:

```
enable icmp address-mask {vlan <vlan name>}

disable icmp address-mask {vlan <vlan name>}
```

To enable or disable the generation of an ICMP parameter-problem message on one or all VLANs, use the following commands:

```
enable icmp parameter-problem {vlan <vlan name>}

disable icmp parameter-problem {vlan <vlan name>}
```

To enable or disable the generation of ICMP port unreachable messages on one or all VLANs, use the following commands:

```
enable icmp port-unreachables {vlan <vlan name>}

disable icmp port-unreachables {vlan <vlan name>}
```

To enable or disable the generation of ICMP redirect messages on one or all VLANs, use the following commands:

```
enable icmp redirects {vlan <vlan name>}

disable icmp redirects {vlan <vlan name>}
```

To enable or disable the generation of ICMP time exceeded messages on one or all VLANs, use the following commands:

```
enable icmp time-exceeded {vlan <vlan name>}
```

```
disable icmp time-exceeded {vlan <vlan name>}
```

To enable or disable the generation of an ICMP timestamp response on one or all VLANs, use the following commands:

```
enable icmp timestamp {vlan <vlan name>}
```

```
disable icmp timestamp {vlan <vlan name>}
```

To enable or disable the generation of ICMP unreachable messages on one or all VLANs, use the following commands:

```
enable icmp unreachables {vlan <vlan name>}
```

```
disable icmp unreachables {vlan <vlan name>}
```

To enable or disable the modification of route table information when an ICMP redirect message is received, use the following commands:

```
enable icmp useredirects
```

```
disable icmp useredirects
```

To reset all of the ICMP settings to the default values, use the following command:

```
unconfigure icmp
```

For ICMP packets that are typically routed, you can apply access lists to restrict forwarding behavior. Access lists are described in Chapter 9.

Configuring DHCP/BOOTP Relay

Once IP unicast routing is configured, you can configure the switch to forward Dynamic Host Configuration Protocol (DHCP) or BOOTP requests coming from clients on subnets being serviced by the switch and going to hosts on different subnets. This feature can be used in various applications, including DHCP services between Windows NT servers and clients running Windows 95. To configure the relay function, follow these steps:

1. Configure VLANs and IP unicast routing.
2. Enable the DHCP or BOOTP relay function, using the following command:

   ```
   enable bootprelay
   ```

3. Configure the addresses to which DHCP or BOOTP requests should be directed, using the following command:

   ```
   configure bootprelay add <ip address>
   ```

To delete a BOOTP relay entry, use the following command:

```
configure bootprelay delete [<ip address> | all]
```

Configuring the DHCP Relay Agent Option (Option 82)

After configuring and enabling the DHCP/BOOTP relay feature, you can enable the DHCP relay agent option feature. This feature inserts a piece of information, called option 82, into any DHCP request

packet that is to be relayed by the switch. Similarly, if a DHCP reply received by the switch contains a valid relay agent option, the option will be stripped from the packet before it is relayed to the client.

The DHCP relay agent option consists of two pieces of data, called sub-options. The first is the agent circuit ID sub-option, and the second is the agent remote ID sub-option. When the DHCP relay agent option is enabled on switches running ExtremeWare, the value of these sub-options is set as follows:

- **Agent circuit ID sub-option**: Contains the ID of the port on which the original DHCP request packet was received. This ID is encoded as (*port_number*). For example, if the DHCP request were received on port 12, the agent circuit ID value would be 3012. On non-slot-based switches, the agent circuit ID value is simply the port number.
- **Agent remote ID sub-option**: Always contains the Ethernet MAC address of the relaying switch. You can display the Ethernet MAC address of the switch by issuing the `show switch` command.

To enable the DHCP relay agent option, use the following command after configuring the DHCP/BOOTP relay function:

```
configure bootprelay dhcp-agent information option
```

To disable the DHCP relay agent option, use the following command:

```
unconfigure bootprelay dhcp-agent information option
```

In some instances, a DHCP server may not properly handle a DHCP request packet containing a relay agent option. To prevent DHCP reply packets with invalid or missing relay agent options from being forwarded to the client, use the following command:

```
configure bootprelay dhcp-agent information check
```

To disable checking of DHCP replies, use this command:

```
unconfigure bootprelay dhcp-agent information check
```

A DHCP relay agent may receive a client DHCP packet that has been forwarded from another relay agent. If this relayed packet already contains a relay agent option, then the switch will handle this packet according to the configured DHCP relay agent option policy. To configure this policy, use the following command:

```
configure bootprelay dhcp-agent information policy <policy>
```

where `<policy>` must be one of the following values: `replace`, `keep`, or `drop`. The default relay policy is `replace`. To configure the policy to the default, use this command:

```
unconfigure bootprelay dhcp-agent information policy
```

For more general information about the DHCP relay agent information option, refer to RFC 3046.

Verifying the DHCP/BOOTP Relay Configuration

To verify the DHCP/BOOTP relay configuration, use the following command:

```
show ipconfig
```

This command displays the configuration of the BOOTP relay service, and the addresses that are currently configured.

UDP-Forwarding

UDP-forwarding is a flexible and generalized routing utility for handling the directed forwarding of broadcast UDP packets. UDP-forwarding allows applications, such as multiple DHCP relay services from differing sets of VLANs, to be directed to different DHCP servers. The following rules apply to UDP broadcast packets handled by this feature:

- If the UDP profile includes BOOTP or DHCP, it is handled according to guidelines in RFC 1542.
- If the UDP profile includes other types of traffic, these packets have the IP destination address modified as configured, and changes are made to the IP and UDP checksums and decrements to the TTL field, as appropriate.

If the UDP-forwarding is used for BOOTP or DHCP forwarding purposes, do not configure or use the existing `bootprelay` function. However, if the previous `bootprelay` functions are adequate, you may continue to use them.

NOTE

UDP-forwarding only works across a layer 3 boundary.

Configuring UDP-Forwarding

To configure UDP-forwarding, the first thing you must do is create a UDP-forward destination profile. The profile describes the types of UDP packets (by port number) that are used, and where they are to be forwarded. You must give the profile a unique name, in the same manner as a VLAN, protocol filter, or Spanning Tree Domain.

Next, configure a VLAN to make use of the UDP-forwarding profile. As a result, all incoming traffic from the VLAN that matches the UDP profile is handled as specified in the UDP-forwarding profile.

A maximum of ten UDP-forwarding profiles can be defined. Each named profile may contain a maximum of eight “rules” defining the UDP port, and destination IP address or VLAN. A VLAN can make use of a single UDP-forwarding profile. UDP packets directed toward a VLAN use an all-ones broadcast on that VLAN.

UDP-Forwarding Example

In this example, the VLAN *Marketing* and the VLAN *Operations* are pointed toward a specific backbone DHCP server (with IP address 10.1.1.1) and a backup server (with IP address 10.1.1.2). Additionally, the VLAN *LabUser* is configured to use any responding DHCP server on a separate VLAN called *LabSvrs*.

The commands for this configuration are as follows:

```
create udp-profile backbonedhcp
create udp-profile labdhcp
configure backbonedhcp add 67 ipaddress 10.1.1.1
configure backbonedhcp add 67 ipaddress 10.1.1.2
configure labdhcp add 67 vlan labsvrs
configure marketing udp-profile backbonedhcp
configure operations udp-profile backbonedhcp
configure labuser udp-profile labdhcp
```

UDP Echo Server

You can use UDP Echo packets to measure the transit time for data between the transmitting and receiving end.

To enable UDP echo server support, use the following command:

```
enable udp-echo-server
```

To disable UDP echo server support, use the following command:

```
disable udp-echo-server
```

13 Interior Gateway Protocols

This chapter describes the following topics:

This chapter assumes that you are already familiar with IP unicast routing. If not, refer to the following publications for additional information:

- RFC 1058—*Routing Information Protocol (RIP)*
- RFC 1723—*RIP Version 2*
- RFC 2178—*OSPF Version 2*
- *Interconnections: Bridges and Routers*
  by Radia Perlman
  ISBN 0-201-56332-0
  Published by Addison-Wesley Publishing Company

Overview

The switch supports the use of two interior gateway protocols (IGPs); the Routing Information Protocol (RIP) and the Open Shortest Path First (OSPF) protocol.

RIP is a distance-vector protocol, based on the Bellman-Ford (or distance-vector) algorithm. The distance-vector algorithm has been in use for many years, and is widely deployed and understood.

OSPF is a link-state protocol, based on the Dijkstra link-state algorithm. OSPF is a newer Interior Gateway Protocol (IGP), and solves a number of problems associated with using RIP on today's complex networks.

NOTE

RIP and OSPF can be enabled on a single VLAN.

RIP Versus OSPF

The distinction between RIP and OSPF lies in the fundamental differences between distance-vector protocols and link-state protocols. Using a distance-vector protocol, each router creates a unique routing table from summarized information obtained from neighboring routers. Using a link-state protocol, every router maintains an identical routing table created from information obtained from all routers in the autonomous system. Each router builds a shortest path tree, using itself as the root. The link-state protocol ensures that updates sent to neighboring routers are acknowledged by the neighbors, verifying that all routers have a consistent network map.

The biggest advantage of using RIP is that it is relatively simple to understand and implement, and it has been the *de facto* routing standard for many years.

RIP has a number of limitations that can cause problems in large networks, including:

- A limit of 15 hops between the source and destination networks.
- A large amount of bandwidth taken up by periodic broadcasts of the entire routing table.
- Slow convergence.
- Routing decisions based on hop count; no concept of link costs or delay.
- Flat networks; no concept of areas or boundaries.

OSPF offers many advantages over RIP, including:

- No limitation on hop count.
- Route updates multicast only when changes occur.
- Faster convergence.
- Support for load balancing to multiple routers based on the actual cost of the link.
- Support for hierarchical topologies where the network is divided into areas.

The details of RIP and OSPF are explained later in this chapter.

Overview of RIP

RIP is an Interior Gateway Protocol (IGP) first used in computer routing in the Advanced Research Projects Agency Network (ARPAnet) as early as 1969. It is primarily intended for use in homogeneous networks of moderate size.

To determine the best path to a distant network, a router using RIP always selects the path that has the least number of hops. Each router that data must traverse is considered to be one hop.

Routing Table

The routing table in a router using RIP contains an entry for every known destination network. Each routing table entry contains the following information:

- IP address of the destination network
- Metric (hop count) to the destination network
- IP address of the next router
- Timer that tracks the amount of time since the entry was last updated

The router exchanges an update message with each neighbor every 30 seconds (default value), or if there is a change to the overall routed topology (also called *triggered updates*). If a router does not receive an update message from its neighbor within the route timeout period (180 seconds by default), the router assumes the connection between it and its neighbor is no longer available.

Split Horizon

Split horizon is a scheme for avoiding problems caused by including routes in updates sent to the router from which the route was learned. Split horizon omits routes learned from a neighbor in updates sent to that neighbor.

To enable split horizon on RIP, issue this command:

```
enable rip splithorizon
```

To disable split horizon on RIP, issue this command:

```
disable rip splithorizon
```

Poison Reverse

Like split horizon, poison reverse is a scheme for eliminating the possibility of loops in the routed topology. In this case, a router advertises a route over the same interface that supplied the route, but the route uses a hop count of 16, defining it as unreachable.

To enable poison reverse, issue this command:

```
enable rip poisonreverse
```

To disable poison reverse, issue this command:

```
disable rip poisonreverse
```

Triggered Updates

Triggered updates occur whenever a router changes the metric for a route, and it is required to send an update message immediately, even if it is not yet time for a regular update message to be sent. This will generally result in faster convergence, but may also result in more RIP-related traffic.

To enable triggered updates on RIP, issue this command:

```
enable rip triggerupdate
```

To disable triggered updates on RIP, issue this command:

```
disable rip triggerupdate
```

Route Advertisement of VLANs

VLANs that are configured with an IP address, but are configured to not route IP or are not configured to run RIP, do not have their subnets advertised by RIP. Only those VLANs that are configured with an IP address and are configured to route IP and run RIP have their subnets advertised.

RIP Version 1 Versus RIP Version 2

A new version of RIP, called RIP version 2, expands the functionality of RIP version 1 to include:

- Variable-Length Subnet Masks (VLSMs).
- Support for next-hop addresses, which allows for optimization of routes in certain environments.
- Multicasting.

  RIP version 2 packets can be multicast instead of being broadcast, reducing the load on hosts that do not support routing protocols.

NOTE

If you are using RIP with supernetting/Classless Inter-Domain Routing (CIDR), you must use RIPv2 only. In addition, RIP route aggregation must be turned off.

Overview of OSPF

OSPF is a link-state protocol that distributes routing information between routers belonging to a single IP domain, also known as an *autonomous system* (AS). In a link-state routing protocol, each router maintains a database describing the topology of the autonomous system. Each participating router has an identical database maintained from the perspective of that router.

From the link-state database (LSDB), each router constructs a tree of shortest paths, using itself as the root. The shortest path tree provides the route to each destination in the autonomous system. When several equal-cost routes to a destination exist, traffic can be distributed among them. The cost of a route is described by a single metric.

Link-State Database

Upon initialization, each router transmits a link-state advertisement (LSA) on each of its interfaces. LSAs are collected by each router and entered into the LSDB of each router. Once all LSAs are received, the router uses the LSDB to calculate the best routes for use in the IP routing table. OSPF uses flooding to distribute LSAs between routers. Any change in routing information is sent to all of the routers in the network. All routers within an area have the exact same LSDB. Table 39 describes LSA type numbers.

Table 39: LSA Type Numbers

Type Number	Description
1	Router LSA
2	Network LSA
3	Summary LSA
4	AS summary LSA
5	AS external LSA
7	NSSA external LSA
9	Link local—Opaque
10	Area scoping—Opaque
11	AS scoping—Opaque

OSPF passive adds the interface to the Type 1 LSA, but it does not send hellos or establish adjacencies on that interface.

Database Overflow

The OSPF database overflow feature allows you to limit the size of the LSDB and to maintain a consistent LSDB across all the routers in the domain, which ensures that all routers have a consistent view of the network.

Consistency is achieved by:

- Limiting the number of external LSAs in the database of each router.
- Ensuring that all routers have identical LSAs.

To configure OSPF database overflow, use the following command:

```
configure ospf ase-limit <number> {timeout <seconds>}
```

where:

- `<number>`—Specifies the number of external LSAs that the system supports before it goes into overflow state. A limit value of zero disables the functionality.

  When the LSDB size limit is reached, OSPF database overflow flushes LSAs from the LSDB. OSPF database overflow flushes the same LSAs from all the routers, which maintains consistency.
- `timeout`—Specifies the timeout, in seconds, after which the system ceases to be in overflow state. A timeout value of zero leaves the system in overflow state until OSPF is disabled and re-enabled.

Opaque LSAs

Opaque LSAs are a generic OSPF mechanism used to carry auxiliary information in the OSPF database. Opaque LSAs are most commonly used to support OSPF traffic engineering.

Normally, support for opaque LSAs is auto-negotiated between OSPF neighbors. In the event that you experience interoperability problems, you can disable opaque LSAs across the entire system using the following command:

```
disable ospf capability opaque-lsa
```

To re-enable opaque LSAs across the entire system, use the following command:

```
enable ospf capability opaque-lsa
```

If your network uses opaque LSAs, we recommend that all routers on your OSPF network support opaque LSAs. Routers that do not support opaque LSAs do not store or flood them. At minimum a well-interconnected subsection of your OSPF network needs to support opaque LSAs to maintain reliability of their transmission.

On an OSPF broadcast network, the designated router (DR) must support opaque LSAs or none of the other routers on that broadcast network will reliably receive them. You can use the OSPF priority feature to give preference to an opaque-capable router, so that it becomes the elected DR.

For transmission to continue reliably across the network, the backup designated router (BDR) must also support opaque LSAs.

NOTE

Opaque LSAs are supported in ExtremeWare version 6.2 and above.

Areas

OSPF allows parts of a network to be grouped together into *areas*. The topology within an area is hidden from the rest of the autonomous system. Hiding this information enables a significant reduction in LSA traffic, and reduces the computations needed to maintain the LSDB. Routing within the area is determined only by the topology of the area.

The three types of routers defined by OSPF are as follows:

- **Internal Router (IR)**—An internal router has all of its interfaces within the same area.
- **Area Border Router (ABR)**—An ABR has interfaces in multiple areas. It is responsible for exchanging summary advertisements with other ABRs. You can create a maximum of 7 non-zero areas.
- **Autonomous System Border Router (ASBR)**—An ASBR acts as a gateway between OSPF and other routing protocols, or other autonomous systems.

Backbone Area (Area 0.0.0.0)

Any OSPF network that contains more than one area is required to have an area configured as area 0.0.0.0, also called the *backbone*. All areas in an autonomous system must be connected to the backbone. When designing networks, you should start with area 0.0.0.0, and then expand into other areas.

NOTE

Area 0.0.0.0 exists by default and cannot be deleted or changed.

The backbone allows summary information to be exchanged between ABRs. Every ABR hears the area summaries from all other ABRs. The ABR then forms a picture of the distance to all networks outside of its area by examining the collected advertisements, and adding in the backbone distance to each advertising router.

When a VLAN is configured to run OSPF, you must configure the area for the VLAN. If you want to configure the VLAN to be part of a different OSPF area, use the following command:

```
configure ospf add vlan area
```

If this is the first instance of the OSPF area being used, you must create the area first using the following command:

```
create ospf area
```

Stub Areas

OSPF allows certain areas to be configured as *stub areas*. A stub area is connected to only one other area. The area that connects to a stub area can be the backbone area. External route information is not distributed into stub areas. Stub areas are used to reduce memory consumption and computation requirements on OSPF routers. Use the following command to configure an OSPF area as a stub area:

```
configure ospf area stub stub-default-cost
```

Not-So-Stubby-Areas (NSSA)

NSSAs are similar to the existing OSPF stub area configuration option, but have the following two additional capabilities:

- External routes originating from an ASBR connected to the NSSA can be advertised within the NSSA.
- External routes originating from the NSSA can be propagated to other areas, including the backbone area.

The CLI command to control the NSSA function is similar to the command used for configuring a stub area, as follows:

```
configure ospf area nssa stub-default-cost
```

The `translate` option determines whether type 7 LSAs are translated into type 5 LSAs. When configuring an OSPF area as an NSSA, the `translate` should only be used on NSSA border routers, where translation is to be enforced. If `translate` is not used on any NSSA border router in a NSSA, one of the ABRs for that NSSA is elected to perform translation (as indicated in the NSSA specification). The option should not be used on NSSA internal routers. Doing so inhibits correct operation of the election algorithm.

Normal Area

A normal area is an area that is not:

- Area 0.

- Stub area.
- NSSA.

Virtual links can be configured through normal areas. External routes can be distributed into normal areas.

Virtual Links

In the situation when a new area is introduced that does not have a direct physical attachment to the backbone, a *virtual link* is used. A virtual link provides a logical path between the ABR of the disconnected area and the ABR of the normal area that connects to the backbone. A virtual link must be established between two ABRs that have a common area, with one ABR connected to the backbone. Figure 44 illustrates a virtual link.

NOTE

Virtual links can not be configured through a stub or NSSA area.

Figure 44: Virtual link using Area 1 as a transit area

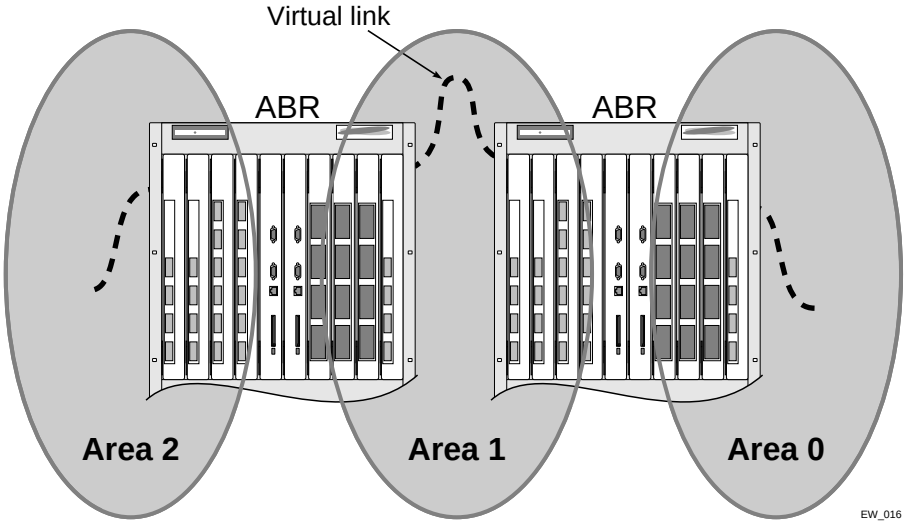

Virtual links are also used to repair a discontiguous backbone area. For example, in Figure 45, if the connection between ABR1 and the backbone fails, the connection using ABR2 provides redundancy so that the discontiguous area can continue to communicate with the backbone using the virtual link.

Figure 45: Virtual link providing redundancy

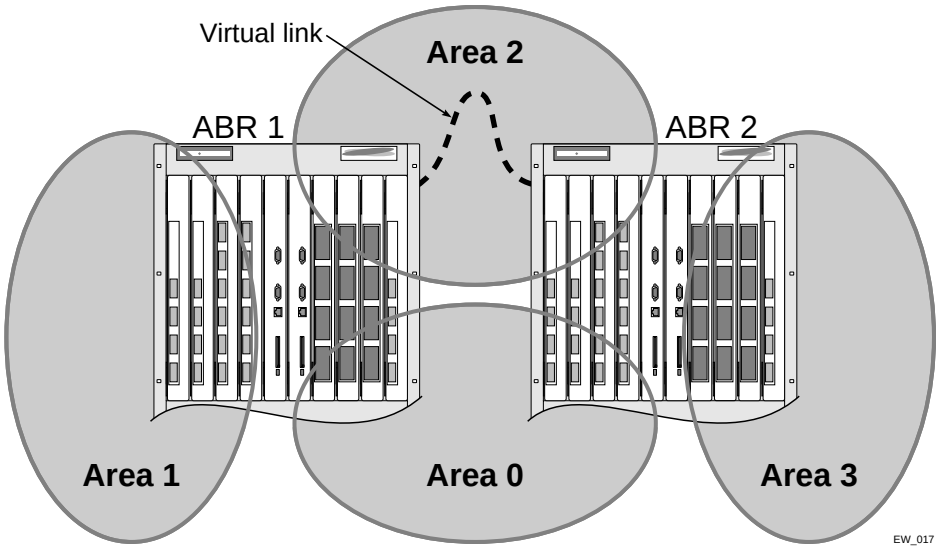

Point-to-Point Support

You can manually configure the OSPF link type for a VLAN. Table 40 describes the link types.

Table 40: OSPF Link Types

Link Type	Number of Routers	Description
Auto	Varies	ExtremeWare automatically determines the OSPF link type based on the interface type. This is the default setting.
Broadcast	Any	Routers must elect a designated router (DR) and a backup designated router (BDR) during synchronization. Ethernet is an example of a broadcast link.
Point-to-point	Up to 2	Synchronizes faster than a broadcast link because routers do not elect a DR or BDR. Does not operate with more than two routers on the same VLAN. PPP is an example of a point-to-point link. An OSPF point-to-point link supports only zero to two OSPF routers and does not elect a DR or BDR. If you have three or more routers on the VLAN, OSPF will fail to synchronize if the neighbor is not configured.
Passive		A passive link does not send or receive OSPF packets.

NOTE

The number of routers in an OSPF point-to-point link is determined per-VLAN, not per-link.

NOTE

All routers in the VLAN must have the same OSPF link type. If there is a mismatch, OSPF attempts to operate, but may not be reliable.

Route Re-Distribution

RIP and OSPF can be enabled simultaneously on the switch. Route re-distribution allows the switch to exchange routes, including static routes, between the three routing protocols. Figure 46 is an example of route re-distribution between an OSPF autonomous system and a RIP autonomous system.

Figure 46: Route re-distribution

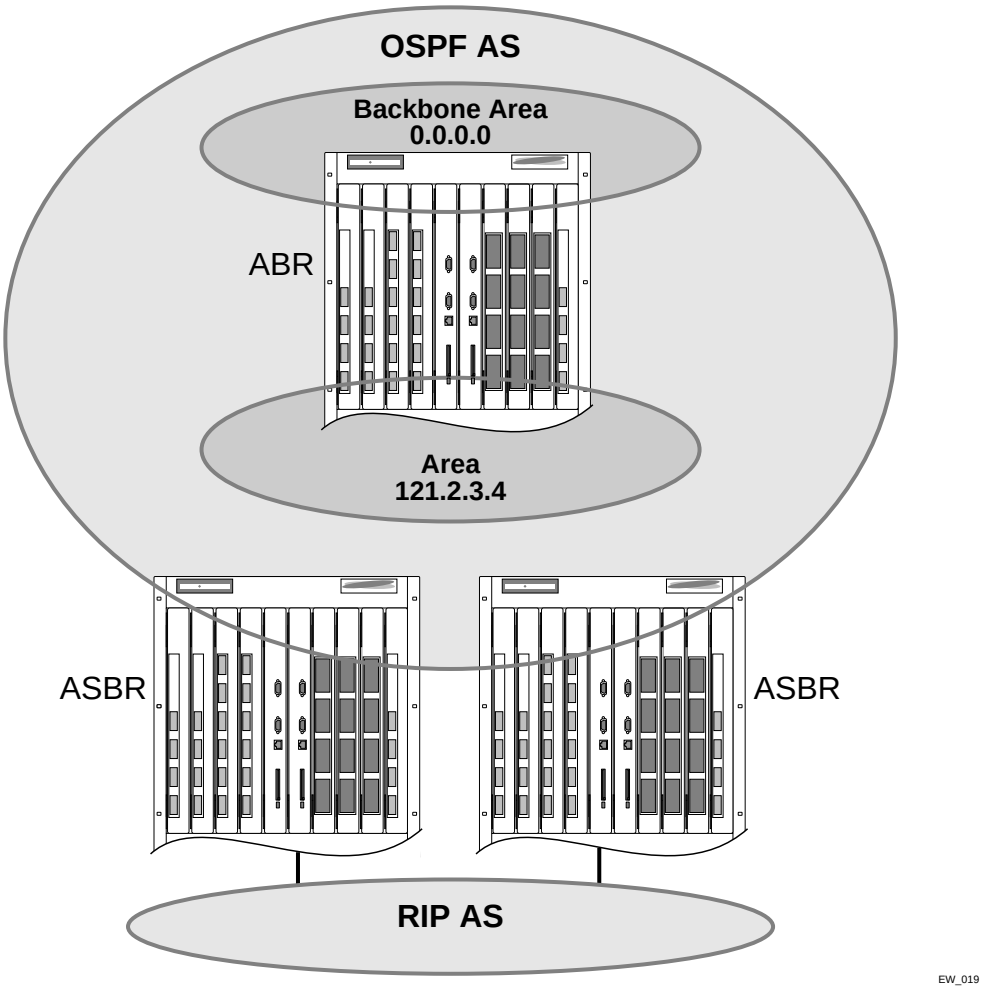

Configuring Route Re-Distribution

Exporting routes from one protocol to another, and from that protocol to the first one, are discreet configuration functions. For example, to run OSPF and RIP simultaneously, you must first configure both protocols and then verify the independent operation of each. Then you can configure the routes to export from OSPF to RIP and the routes to export from RIP to OSPF. Likewise, for any other combinations of protocols, you must separately configure each to export routes to the other.

Re-Distributing Routes into OSPF

Enable or disable the exporting of RIP, static, and direct (interface) routes to OSPF using the following commands:

```
enable ospf export [direct | rip | static] [cost <number> [ase-type-1 | ase-type-2]
{tag <number>}]
```

These commands enable or disable the exporting of RIP, static, and direct routes by way of LSA to other OSPF routers as AS-external type 1 or type 2 routes. The default setting is disabled.

The cost metric is inserted for all RIP, static, and direct routes injected into OSPF. If the cost metric is set to 0, the cost is inserted from the route. The tag value is used only by special routing applications. Use 0 if you do not have specific requirements for using a tag. The tag value in this instance has no relationship with 802.1Q VLAN tagging.

The same cost, type, and tag values can be inserted for all the export routes, or route maps can be used for selective insertion. When a route map is associated with the export command, the route map is applied on every exported route. The exported routes can also be filtered using route maps. Routes filtered with a route map will be exported as ase-type-1.

Enable or disable the export of virtual IP addresses to other OSPF routers using the following commands:

```
enable ospf export [direct | rip | static] [cost <number> [ase-type-1 | ase-type-2]
{tag <number>}]

disable ospf export [direct | rip | static]
```

Verify the configuration using the command:

```
show ospf
```

Previous Release Issues with OSPF Re-Distribution

In versions of ExtremeWare prior to release 6.0, direct routes corresponding to the interfaces on which RIP was enabled were exported into OSPF as part of RIP routes, using the command `enable ospf export rip`. Using ExtremeWare 6.0 and above, you must configure ExtremeWare to export these direct routes to OSPF. You can use an access profile to filter unnecessary direct routes, using the command:

```
configure ospf direct-filter
```

Re-Distributing Routes into RIP

Enable or disable the exporting of static, direct, and OSPF-learned routes into the RIP domain using the following commands:

```
disable rip export [direct | | ospf | ospf-extern1 | ospf-extern2 | ospf-inter |
ospf-intra | static]

disable rip export [direct | | ospf | ospf-extern1 | ospf-extern2 | ospf-inter |
ospf-intra | static]
```

These commands enable or disable the exporting of static, direct, and OSPF-learned routes into the RIP domain. You can choose which types of OSPF routes are injected, or you can simply choose `ospf`, which will inject all learned OSPF routes regardless of type. The default setting is disabled.

RIP Configuration Example

A switch that has three VLANs is defined as follows:

- *Finance*
  - Contain ports 5 and 6.
  - IP address 192.207.35.1.
- *Personnel*
  - Contain ports 22 and 23.
  - IP address 192.207.36.1.

In this configuration, all IP traffic from stations connected to ports 5 and 6 have access to the switch by way of the VLAN *Finance*. Ports 22 and 23 reach the switch by way of the VLAN *Personnel*.

The example is configured as follows:

```
create vlan Finance
create vlan Personnel

configure Finance protocol ip
configure Personnel protocol ip

configure Finance add port 5, 6
configure Personnel add port 22, 23

configure Finance ipaddress 192.207.35.1
configure Personnel ipaddress 192.207.36.1

enable ipforwarding
configure rip add vlan all
enable rip
```

Configuring OSPF

Each switch that is configured to run OSPF must have a unique router ID. It is recommended that you manually set the router ID of the switches participating in OSPF, instead of having the switch automatically choose its router ID based on the highest interface IP address. Not performing this configuration in larger, dynamic environments could result in an older link state database remaining in use.

Configuring OSPF Wait Interval

ExtremeWare allows you to configure the OSPF wait interval, rather than using the router dead interval.

CAUTION

Do not configure OSPF timers unless you are comfortable exceeding OSPF specifications. Non-standard settings might not be reliable under all circumstances.

To specify the timer intervals, use the following command:

```
configure ospf vlan <vlan name> timer <retransmit interval>
```

You can configure the following parameters:

- **Retransmit interval**—The length of time that the router waits before retransmitting an LSA that is not acknowledged. If you set an interval that is too short, unnecessary retransmissions will result. The default value is 5 seconds.

NOTE

The OSPF standard specifies that wait times are equal to the dead router wait interval.

OSPF Configuration Example

Figure 47 is an example of an autonomous system using OSPF routers. The details of this network follow.

Figure 47: OSPF configuration example

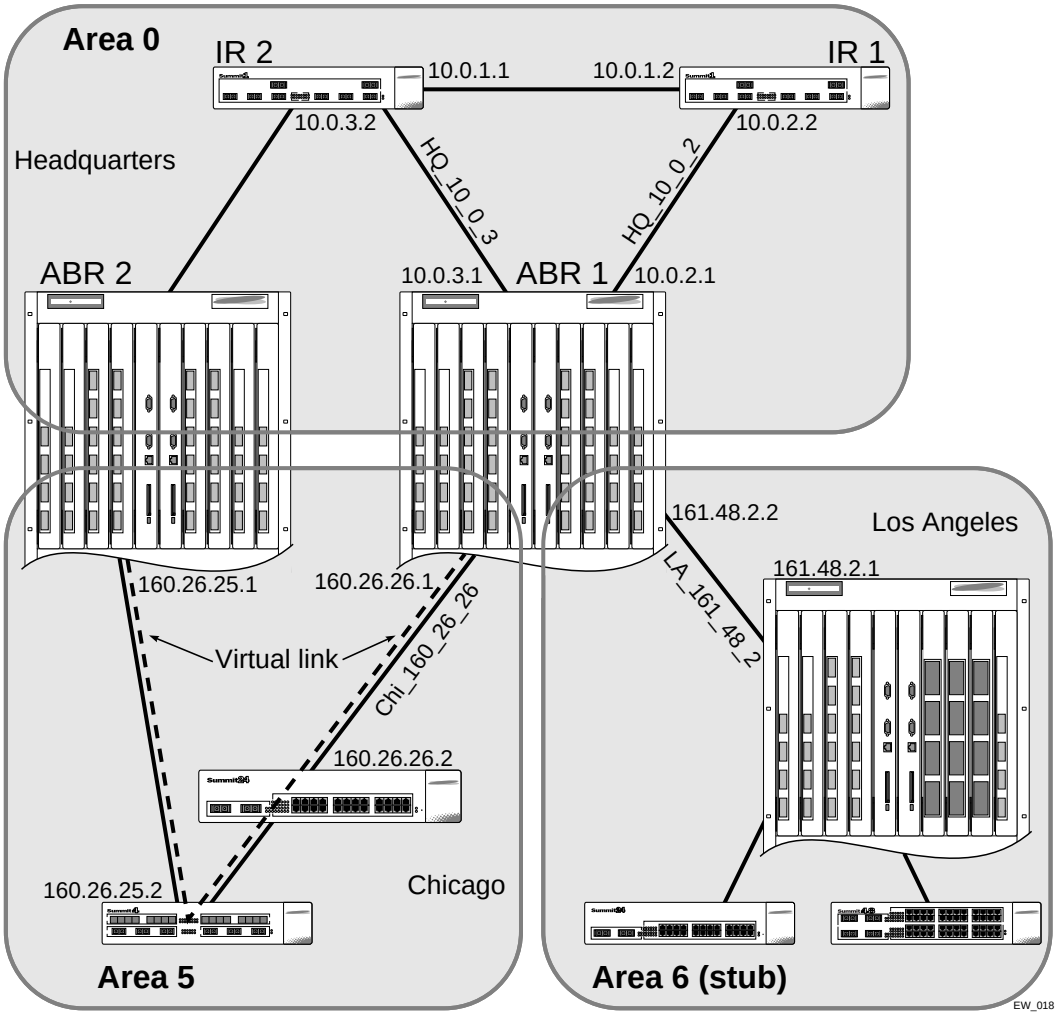

Area 0 is the backbone area. It is located at the headquarters and has the following characteristics:

- Two internal routers (IR1 and IR2)
- Two area border routers (ABR1 and ABR2)
- Network number 10.0.x.x
- Two identified VLANs (HQ_10_0_2 and HQ_10_0_3)

Area 5 is connected to the backbone area by way of ABR1 and ABR2. It is located in Chicago and has the following characteristics:

- Network number 160.26.x.x
- One identified VLAN (Chi_160_26_26)
- Two internal routers

Area 6 is a stub area connected to the backbone by way of ABR1. It is located in Los Angeles and has the following characteristics:

- Network number 161.48.x.x
- One identified VLAN (LA_161_48_2)
- Three internal routers
- Uses default routes for inter-area routing

Two router configurations for the example in Figure 47 are provided in the following section.

Configuration for ABR1

The router labeled ABR1 has the following configuration:

```
create vlan HQ_10_0_2
create vlan HQ_10_0_3
create vlan LA_161_48_2
create vlan Chi_160_26_26

configure vlan HQ_10_0_2 ipaddress 10.0.2.1 255.255.255.0
configure vlan HQ_10_0_3 ipaddress 10.0.3.1 255.255.255.0
configure vlan LA_161_48_26 ipaddress 161.48.2.26 255.255.255.0
configure vlan Chi_160_26_26 ipaddress 160.26.2.1 255.255.255.0

create ospf area 0.0.0.5
create ospf area 0.0.0.6

enable ipforwarding

configure ospf area 0.0.0.6 stub nosummary stub-default-cost 10
configure ospf add vlan LA_161_48_2 area 0.0.0.6
configure ospf add vlan Chi_160_26_26 area 0.0.0.5
configure ospf add vlan all area 0.0.0.0

enable ospf
```

Configuration for IR1

The router labeled IR1 has the following configuration:

```
configure vlan HQ_10_0_1 ipaddress 10.0.1.2 255.255.255.0
configure vlan HQ_10_0_2 ipaddress 10.0.2.2 255.255.255.0
enable ipforwarding
configure ospf add vlan all area 0.0.0.0
enable ospf
```

Displaying OSPF Settings

There are a number of commands you can use to display settings for OSPF. To show global OSPF information, use the `show ospf` command with no options.

To display information about one or all OSPF areas, use the following command:

```
show ospf area <area identifier>
```

The `detail` option displays information about all OSPF areas in a detail format.

To display information about OSPF interfaces for an area, a VLAN, or for all interfaces, use the following command:

```
show ospf interfaces {vlan <vlan name> | area <area identifier>}
```

The `detail` option displays information about all OSPF interfaces in a detail format.

OSPF LSDB Display

ExtremeWare provides several filtering criteria for the `show ospf lsdb` command. You can specify multiple search criteria and only results matching all of the criteria are displayed. This allows you to control the displayed entries in large routing tables.

To display the current link-state database, use the following command:

```
show ospf lsdb area [all | <area identifier>[/<len>] | detail | interface | lsid
<id>[/<len>] | lstype [all | as-external | external-type7 | network | opaque-area |
opaque-global | opaque-local | router | summary-asb |summary-net| routerid
<id>[/<len>] | stats | summary | vlan <vlan name>]
```

The `detail` option displays all fields of matching LSAs in a multi-line format. The `summary` option displays several important fields of matching LSAs, one line per LSA. The `stats` option displays the number of matching LSAs, but not any of their contents. If not specified, the default is to display in the summary format.

A common use of this command is to omit all optional parameters, resulting in the following shortened form:

```
show ospf lsdb
```

The shortened form displays all areas and all types in a summary format.

Authentication

Authentication is supported at two different levels: interface, and domain or area.

- **Interface authentication**—prevents unauthorized routers from forming adjacency. This is achieved by inserting authentication information in the Hello PDUs and validating them on the received Hello PDUs. You can configure authentication separately for level 1 and level 2.
- **Domain or area authentication**—prevents intruders from injecting invalid routing information into this router. Similar to interface authentication, this is achieved by inserting the authentication information using LSP, CSNP, and PSNP PDUs and validating them on receipt. You can configure authentication separately for level 1 and level 2.

At each of the above levels two different authentication methods are supported: simple password as specified in ISO/IEC 10589, and HMAC-MD5 as specified in draft-ietf-isis-hmac-00.txt.

Summarizing Level 1 IP Routing Information

Level 2 routers include in their level 2 LSPs a list of all combinations (IP address, subnet mask, and metric) reachable in the level 1 area attached to them. This information is gathered from the level 1 LSPs from all routers in the area. By default the combinations from all the level 1 routers are included in the level 2 LSPs. Summarization of the level 1 combinations reduces the amount of information stored on the level 2 router and helps in scaling to a large routing domain.

You can configure the level 1 areas with one or more combinations for announcement in their level 2 LSPs. The level 1 IP routing information is matched against the summary addresses configured on the level 1 area. Matches are included in the level 2 LSP.

You can also configure the level 2 router to disregard the summary information. This effectively acts as a filter, preventing reachability information from being included in the level 2 LSP.

Filtering Level 1 IP Routing Information

Level 2 routers include in their level 2 LSPs a list of all combinations (IP address, subnet mask, and metric) reachable in the level 1 area attached to them. This information is gathered from the level 1 LSPs from all routers in the area. By default the combinations from all the level 1 routers are included in the level 2 LSPs. Filtering the level 1 combinations prevents the advertisement of the information to other parts of the domain. This creates a network that is reachable only from routers within the area.

You can configure the level 1 areas in the router with an IP access profile. The level 1 IP routing information in the level 2 LSP is matched against the access profile, and if the result is a deny, the information is not included in the level 2 LSP.

Originating Default Route

This feature injects IP routing information for the default route in the LSP originated by the router, thereby advertising the router as the default gateway.

Injection of the default route into the level 2 subdomain and level 1 area can be controlled individually. You can configure the metric and metric type associated with the default route. You can also configure the default to be automatically generated based on the presence of a default route in the kernel routing table.

Overload Bit

This feature forces the router to set the overload bit (also known as the hippity bit) in its non-pseudo node link-state packets. Normally the setting of the overload bit is allowed only when a router runs into

problems. For example, when a router has a memory shortage, it might be that the Link State database is not complete, resulting in an incomplete or inaccurate routing table. By setting the overload bit in its LSPs, other routers can ignore the unreliable router in their SPF calculations until the router has recovered from its problems.

Set the overload bit when you want to prevent traffic flow.

Default Routes to Nearest Level 1/2 Switch for Level 1 Only Switches

When one router is a level 1 switch, the route to the nearest level 1/2 switch which attaches to a level 2 backbone network may be installed in the kernel routing table of the level 1 switch.

There are three kinds of level 1 only switches:

- a switch that does not attach to any level 1/2 switch; it is part of a level 1 only network
- a switch that attaches to at least one level 1/2 switch, but none of the level 1/2 switches are attached to a level 2 backbone network. Here the level 1 non-pseudo node LSP of the level 1/2 switches should set the attach bit to 0. A level 1 only switch will not install the default routes based on the unattached level 1/2 switch's LSP information.
- a switch that attaches to at least one level 1/2 switch, and at least one of the level 1/2 switches is attached to the level 2 backbone network. Here the level 1 non-pseudo node LSP of the level 1/2 switch should set the attach bit to 1. A level 1 only switch will install the default routes based on the attached level 1/2 switch's LSP information.

The level 1/2 switch that is attached to the level 2 backbone network when at least one of area addresses of level 2 LSP received from other level 2 or level 1/2 switches is not in the list of the level 1 union area address set.

14 IP Multicast Routing

This chapter covers the following topics:

For more information on IP multicasting, refer to the following publications:

- RFC 1112 – *Host Extension for IP Multicasting*
- RFC 2236 – *Internet Group Management Protocol, Version 2*
- PIM-SM Version 2 – *draft_ietf_pim_sm_v2_new_04*

The following URLs point to the Web sites for the IETF Working Groups:

IEFT PIM Working Group:

http://www.ietf.org/html.charters/pim-charter.html

IP Multicast Routing Overview

IP multicast routing is a function that allows a single IP host to send a packet to a group of IP hosts. This group of hosts can include devices that reside on the local network, within a private network, or outside of the local network.

IP multicast routing consists of the following functions:

- A router that can forward IP multicast packets.
- A router-to-router multicast routing protocol (such as Protocol Independent Multicast- Sparse Mode (PIM-SM).
- A method for the IP host to communicate its multicast group membership to a router (for example, Internet Group Management Protocol (IGMP)).

NOTE

You should configure IP unicast routing before you configure IP multicast routing.

PIM Sparse Mode (PIM-SM) Overview

Protocol independent Multicast-Sparse Mode (PIM-SM) routes multicast packets to multicast groups. The sparse mode protocol is designed for installations where the multicast groups are scattered over a large area such as a wide area network (WAN). PIM-SM is a router-to-router protocol, so all routers and switches must upgrade to the same PIM-SM version. Summit 400 switches use PIM-SM version 2 to forward IP packets that are destined to the IP addresses in the Class D Range to multiple networks using the Multicast Routing information setup.

PIM-SM is an explicit join and prune protocol that is a mixture of the shared tree and shortest path tree (SPT) models. The routers must explicitly join the group(s) in which they are interested in becoming a member, which is beneficial for large networks that have group members who are sparsely distributed. PIM-SM is not dependant on a specific unicast routing protocol. The Summit 400 supports IGMP, which allows network hosts to report the multicast group membership to the switch.

Using PIM-SM, the source router sends a join message to a known rendezvous point (RP). The RP is a central multicast router that is responsible for receiving and distributing multicast packets. RPs are elected by a bootstrap router (BSR). The job of the BSR is to broadcast bootstrap messages, disseminate RP information, and to elect the RP. You may only configure the Summit 200 switches as an RP in static mode, which means that all switches in your network must be configured with the same RP address for the same group (range). Summit 400 switches are not eligible to be BSRs.

When a source router has a multicast packet to distribute, it encapsulates the packet in a unicast message and sends it to the RP. The RP decapsulates the multicast packet and distributes it among all member routers.

When a router determines that the multicast rate has exceeded a configured threshold, that router can send an explicit join to the originating router. Once this occurs, the receiving router gets the multicast directly from the sending router, and bypasses the RP.

Configuring PIM-SM

You can configure two active and 254 passive interfaces on a Summit 400 for PIM-SM. By default the interface is configured as active. To enable the interface as passive, specify the `passive` keyword; to enable the interface as active, omit the `passive` keyword. The following command enables PIM-SM on an IP interface.

```
configure pim add {vlan} [<vlan name>]
```

The following command disables PIM-SM on an IP interface:

```
configure pim delete vlan [<vlan name> | all]
```

For example, to add a VLAN named *lobby*, as an active interface, you would enter:

```
configure pim add vlan lobby
```

To configure an RP and its associated groups statically, enter the following command:

```
configure pim crp static <rp address> [none | <access profile>] {<priority [0-254]>}
```

The `access profile` contains a list of multicast group accesses served by the RP.

For example, the following command statically configures an RP and its associated groups defined in access profile *rp-list*:

```
configure pim crp static 10.0.3.1 rp-list
```

To configure the candidate RP advertising interval for PIM-SM timers, enter this command:

```
configure pim timer <hello interval> <join prune interval> vlan [<vlan name>]
```

Specify the intervals in seconds. The `hello interval` specifies the amount of time before a hello message is sent out by the PIM router. The `join prune` interval is the amount of time before a join or a prune command is executed. The valid range for both intervals is 1 to 65,519 seconds. The default for the `hello interval` is 30 seconds; the default for `join prune` is 60 seconds.

Because PIM leverages the unicast routing capability that is already present in the switch, the access policy capabilities are, by nature, different. When the PIM protocol is used for routing IP multicast traffic, the switch can be configured to use an access profile to determine trusted PIM router neighbors for the VLAN on the switch running PIM. To configure a trusted neighbor policy enter the following command:

```
configure pim vlan [<vlan name> | all] trusted-gateway [<access profile> | none]
```

For example, the following command configures a trusted neighbor policy on the VLAN *backbone*:

```
configure pim vlan backbone trusted-gateway
```

To configure the threshold (in Kbps) for switching to SPT, enter the following command:

```
configure pim spt-threshold <last hop router threshold> {<rp threshold>}
```

On leaf routers, this setting is based on data packets. On the RP, this setting is based on register packet rate in Kbps.

The following command configures the checksum computation to either include data (for compatibility with Cisco Systems products) or to exclude data (for RFC-compliant operation), in the register message:

```
configure pim register-checksum-to [include-data | exclude-data]
```

IGMP Overview

IGMP is a protocol used by an IP host to register its IP multicast group membership with a router. Periodically, the router queries the multicast group to see if the group is still in use. If the group is still active, a single IP host responds to the query, and group registration is maintained.

IGMP is enabled by default on the switch. However, the switch can be configured to disable the generation of periodic IGMP query packets. IGMP should be enabled when the switch is configured to perform IP unicast or IP multicast routing.

IGMP Snooping

IGMP snooping is a layer 2 function of the switch. It does not require multicast routing to be enabled. In IGMP snooping, the layer 2 switch keeps track of IGMP requests, and only forwards multicast traffic to the part of the local network that requires it. IGMP snooping optimizes the usage of network bandwidth, and prevents multicast traffic from being flooded to parts of the local network that do not need it. The switch does not reduce any IP multicast traffic in the local multicast domain (224.0.0.x).

IGMP snooping is enabled by default on the switch. If IGMP snooping is disabled, all IGMP and IP multicast traffic floods within a given VLAN. IGMP snooping expects at least one device on every VLAN to periodically generate IGMP query messages. The static IGMP snooping entries do not require periodic query. An optional optimization for IGMP snooping is the strict recognition of multicast routers only if the remote devices have joined the PIM (244.0.0.13) multicast groups.

When a port sends an IGMP leave message, the switch removes the IGMP snooping entry after 1000 milli-seconds (the leave time is configurable, ranging from 0 to 10000 ms). The switch sends a query to determine which ports want to remain in the multicast group. If other members of the VLAN want to remain in the multicast group, the router ignores the leave message, but the port that requests removal is removed from the IGMP snooping table.

If the last port within a VLAN sends an IGMP leave message, then the router will not receive any responses to the query, and the router immediately will remove the VLAN from the multicast group.

Static IGMP

In order to receive multicast traffic, a host needs to explicitly join a multicast group by sending an IGMP request, then the traffic is forwarded to that host. There are situations where you would like multicast traffic to be forwarded to a port where a multicast enabled host is not available (for example, testing multicast configurations). Static IGMP emulates a host or router attached to a switch port, so that multicast traffic will be forwarded to that port. Emulate a host to forward a particular multicast group to a port; emulate a router to forward all multicast groups to a port. Use the following command to emulate a host on a port:

```
configure igmp snooping vlan <vlan name> ports <portlist> add static group <ip
address>
```

Use the following command to emulate a multicast router on a port:

```
configure igmp snooping vlan <vlan name> ports <portlist> add static router
```

To remove these entries, use the corresponding command:

```
configure igmp snooping vlan <vlan name> ports <portlist> delete static group [<ip
address> | all]
configure igmp snooping vlan <vlan name> ports <portlist> delete static router
```

To display the IGMP snooping static groups, use the following command:

```
show igmp snooping {vlan <vlan name>} static group
```

IGMP Snooping Filters

IGMP snooping filters allow you to configure an access profile on a port to allow or deny IGMP report and leave packets coming into the port. For details on creating access profiles, see the section, “Routing Access Profiles” on page 156. For the access profiles used as IGMP snooping filters, all the profile entries

should IP address type entries, and the IP address of each entry must be in the class-D multicast address space, but should not be in the multicast control subnet range (224.0.0.x/24). After you have created an access profile, use the following command to associate the access profile and filter with a set of ports:

```
configure igmp snooping vlan <vlan name> ports <portlist> filter [<access profile> | none]
```

To remove the filter, use the none option as shown in the following example:

```
configure igmp snooping vlan <vlan name> ports <portlist> filter none
```

To display the IGMP snooping filters, use the following command:

```
show igmp snooping {vlan <vlan name>} filter
```

Multicast Tools

ExtremeWare provides two commonly available tools to monitor and troubleshoot IP multicast, `mrinfo` and `mtrace`.

Mrinfo

The multicast router information tool, (mrinfo), requests information from a router that could be used for tracing and troubleshooting. A request is sent to a multicast router, and the router responds with the following information:

- code version
- system multicast information
- interface information
  - interface IP address
  - interface multicast capabilities
  - metric configured on the interface
  - threshold configured on the interface
  - count and IP address of the neighbors discovered on the interface

Use the following command to send an `mrinfo` request:

```
mrinfo <ip address> {from <ip address>} {timeout <seconds>}
```

Mtrace

Multicast trace (mtrace) relies on a feature of multicast routers that is accessed using the IGMP protocol. Since multicast uses reverse path forwarding, a multicast trace is run from the destination to the source. A query packet is sent to the last-hop multicast router. This router builds a trace response packet, fills in a report for its hop, and forwards the packet to the next upstream router. As the request is forwarded, each router in turn adds its own report to the trace response. When the request reaches the first-hop router, the filled in request is sent back to the system requesting the trace. The request will also be returned if the maximum hop limit is reached.

If a router does not support the mtrace functionality, it will silently drop the request packet and no information will be returned. For this situation, you could send the trace with a small number of maximum hops allowed, increasing the number of hops as the stream is traced.

The group IP address must be in the class-D multicast address space, but should not be in the multicast control subnet range (224.0.0.x / 24).

Use the following command to trace a multicast stream:

```
mtrace source <ip address> {destination <ip address>} {group <ip address>} {from <ip address>} {gateway <ip address >} {timeout <seconds>} {maximum-hops <number>}
```

Configuring IP Multicasting Routing

To configure IP multicast routing, you must do the following:

1. Configure the system for IP unicast routing.
2. Enable multicast routing on the interface using the following command:

   ```
   enable ipmcforwarding {vlan <vlan name>}
   ```
3. Enable PIM on all IP multicast routing interfaces using the following command:

   ```
   configure pim add {vlan} [<vlan name>]
   ```
4. Enable PIM on the router using one of the following commands:

   ```
   enable pim
   ```

Configuration for IR1

The router labeled IR1 has the following configuration:

```
configure vlan HQ_10_0_1 ipaddress 10.0.1.2 255.255.255.0
configure vlan HQ_10_0_2 ipaddress 10.0.2.2 255.255.255.0
configure ospf add vlan all
enable ipforwarding
enable ospf
enable ipmcforwarding
configure pim add vlan HQ_10_0_1
enable pim
```

The following example configures PIM-SM.

Figure 48: IP multicast routing using PIM-SM configuration example

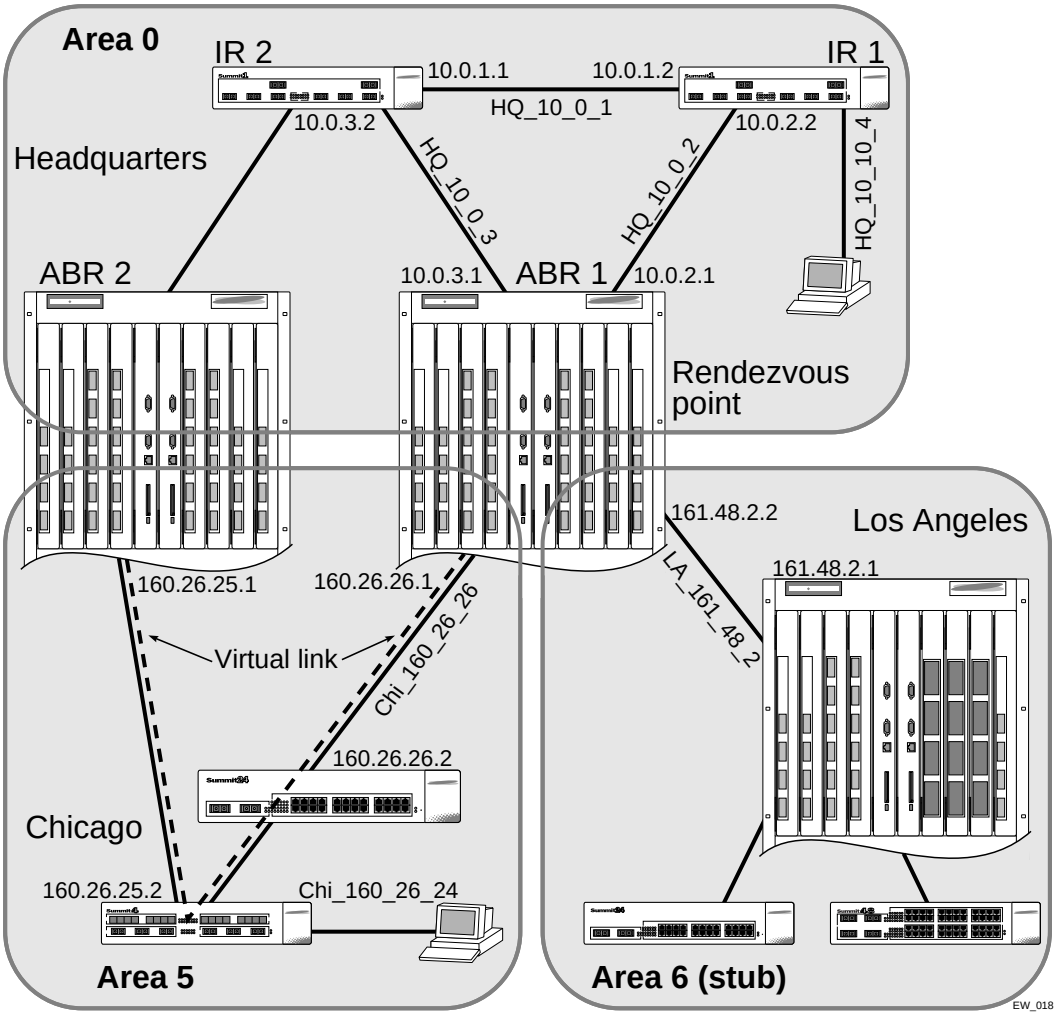

Configuration for ABR1

The router labeled ABR1 has the following configuration:

```
configure vlan HQ_10_0_2 ipaddress 10.0.2.1 255.255.255.0
configure vlan HQ_10_0_3 ipaddress 10.0.3.1 255.255.255.0
configure vlan LA_161_48_2 ipaddress 161.48.2.2 255.255.255.0
configure vlan CHI_160_26_26 ipaddress 160.26.26.1 255.255.255.0
configure ospf add vlan all
enable ipforwarding
enable ipmcforwarding
configure pim add vlan all sparse
create access-profile rp-list ipaddress
configure rp-list add ipaddress 224.0.0.0 240.0.0.0
enable loopback HQ_10_0_3
configure pim crp HQ_10_0_3 rp-list 30
configure pim cbsr HQ_10_0_3 30

configure pim spt-threshold 16 8
```

15 Using ExtremeWare Vista on the Summit 400

This chapter describes the following topics:

ExtremeWare Vista Overview

A standard device-management feature on the Summit 400 is ExtremeWare Vista. Using a web browser, ExtremeWare Vista allows you to access the switch over a TCP/IP network. ExtremeWare Vista provides a subset of the command-line interface (CLI) in a graphical format that allows you to configure the switch and review statistical reports. However because ExtremeWare Vista includes only a subset of the CLI, some commands for the Summit 400 are not available using ExtremeWare Vista. If a particular command is not represented in ExtremeWare Vista, you must use the CLI to achieve the desired result.

Before attempting to access ExtremeWare Vista, ensure:

- You assign an IP address to a VLAN to access the switch. For more information on assigning an IP address, see “Configuring Switch IP Parameters” on page 60.
- You have a properly configured standard web browser that supports frames and JavaScript (such as Netscape Navigator 3.0 or above, or Microsoft Internet Explorer 3.0 or above).

Setting Up Your Browser

In general, the default settings that come configured on your browser work well with ExtremeWare Vista. The following are recommended settings that you can use to improve the display features and functions of ExtremeWare Vista:

- After downloading a newer version of the switch image, clear the browser disk and memory cache to see the updated menus. You must clear the cache while on the main ExtremeWare Vista Logon page, so that all underlying.GIF files are updated.

- Check for newer versions of stored pages. Every visit to the page should be selected as a cache setting.

  If you are using Netscape Navigator, configure the cache option to check for changes “Every Time” you request a page.

  If you are using Microsoft Internet Explorer, configure the Temporary Internet Files setting to check for newer versions of stored pages by selecting “Every visit to the page.”
- On older-browsers you might need to specify that images be auto-loaded.
- Use a high-resolution monitor to maximize the amount of information displayed in the content frame. The recommended resolution is 1024 x 768 pixels. You can also use 800 x 600 pixels.
- Turn off one or more of the browser toolbars to maximize the viewing space of the ExtremeWare Vista content screen.
- If you will be using ExtremeWare Vista to send an email to the Extreme Networks Technical Support department, configure the email settings in your browser.
- Configure the browser to use the following recommended fonts:

  — Proportional font—Times New Roman

  — Fixed-width font—Courier New

Accessing ExtremeWare Vista

After an IP address is assigned to the VLAN, you can access the default home page of the switch.

1 Enter the following command in your browser:

http://<ipaddress>

The home page for the Summit 400 opens as shown in Figure 49.

Figure 49: Home Page for ExtremeWare Vista

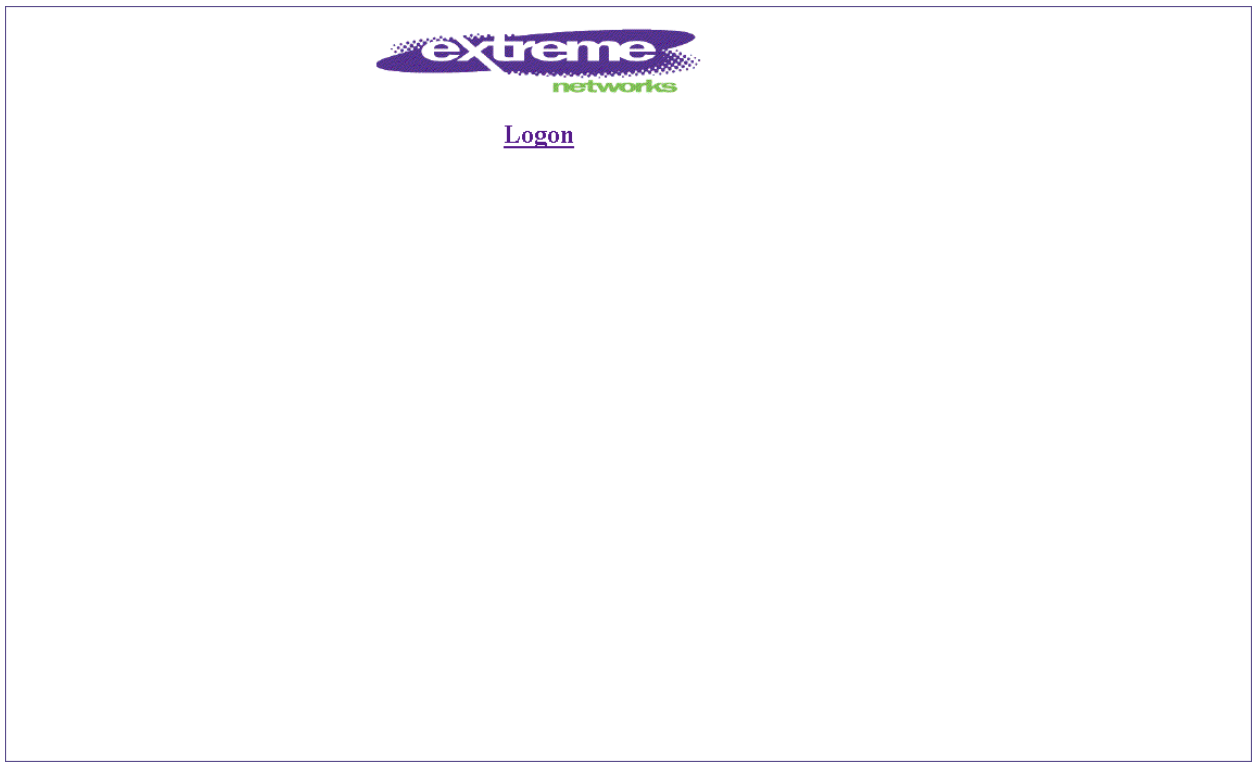

2 Click Logon to open the Username and Password dialog box shown in Figure 50.

Figure 50: Username and Password Dialog Box

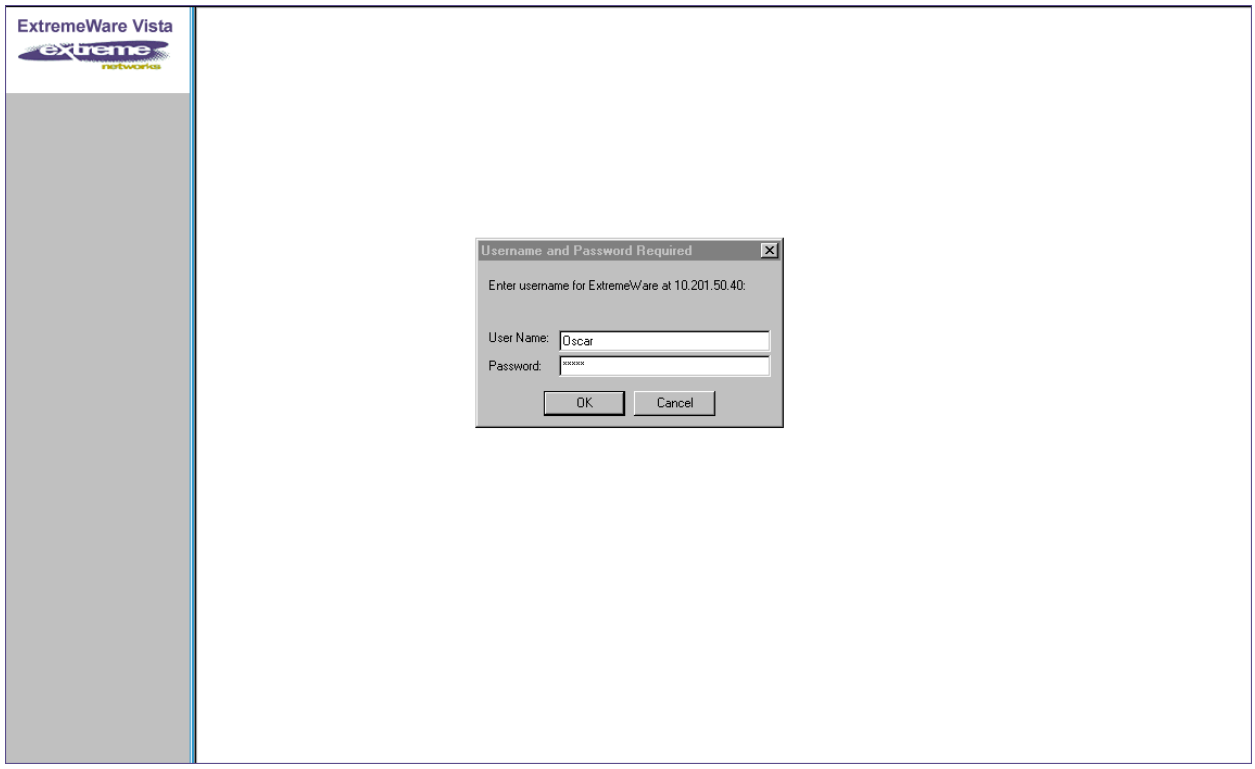

3 Type your username and password and click **OK**. The main page for the switch opens as shown in Figure 51.

If you enter the username and password of an administrator-level account, you have access to all ExtremeWare Vista pages. If you enter a user-level account name and password, you only have access to the Statistics and Support information.

Figure 51: Summit 400 Main Page

Navigating within ExtremeWare Vista

ExtremeWare Vista pages use a common HTML frameset comprised of two frames: a content frame and a task frame. The content frame contains the main body of information in ExtremeWare Vista. The task frame contains a menu of four buttons that correspond to the four main functions:

- Configuration
- Statistics
- Support
- Logout

While these buttons can be expanded or contracted to display the submenu links, all four main functions are static in that they are visible at all times during the session.

When you choose one of the main buttons, that menu expands to reveal the submenu links available under that function. If another function list is open at the time, that list contracts so that only the active menu is open.

When you choose a submenu link in the task frame, the content frame populates with the corresponding data. However when you choose a new task, the content frame does not change until you choose a new a submenu link and repopulate the frame.

Browser Controls

Browser controls include drop-down list boxes, check boxes, and multiselect list boxes. A multiselect list box has a scrollbar on the right side of the box. Using a multiselect list box, you can select a single item, all items, a set of contiguous items, or multiple noncontiguous items. Table 41 describes how to make selections from a multiselect list box.

Table 41: Multiselect List Box Key Definitions

Selection Type	Key Sequence
Single item	Click the item using the mouse.
All items	Click the first item, and drag to the last item.
Contiguous items	Click the first desired item, and drag to the last desired item.
Selected noncontiguous items	Hold down [Ctrl], click the first desired item, click the next desired item, and so on.

Status Messages

Status messages are displayed at the top of the content frame. The four types of status messages are:

- **Information**—Displays information that is useful to know before, or as a result of, changing configuration options.
- **Warning**—Displays warnings about the switch configuration.
- **Error**—Displays errors caused by incorrectly configured settings.
- **Success**—Displays informational messages after you click Submit. The message displayed reads, “Request was submitted successfully.” These informational messages indicate that the operation was successful.

Standalone Buttons

At the bottom of some of the content frames is a section that contains standalone buttons. Standalone buttons are used to perform tasks that are not associated with a particular configuration option. An example of this is the Reboot Switch button.

Configuring the Summit 400 using ExtremeWare Vista

You can configure many features of the Summit 400-48. Click the Configuration button in the task frame to reveal the submenu links, as shown in Figure 52. These configuration tasks are described in the following sections:

Figure 52: Configuration Submenu Links

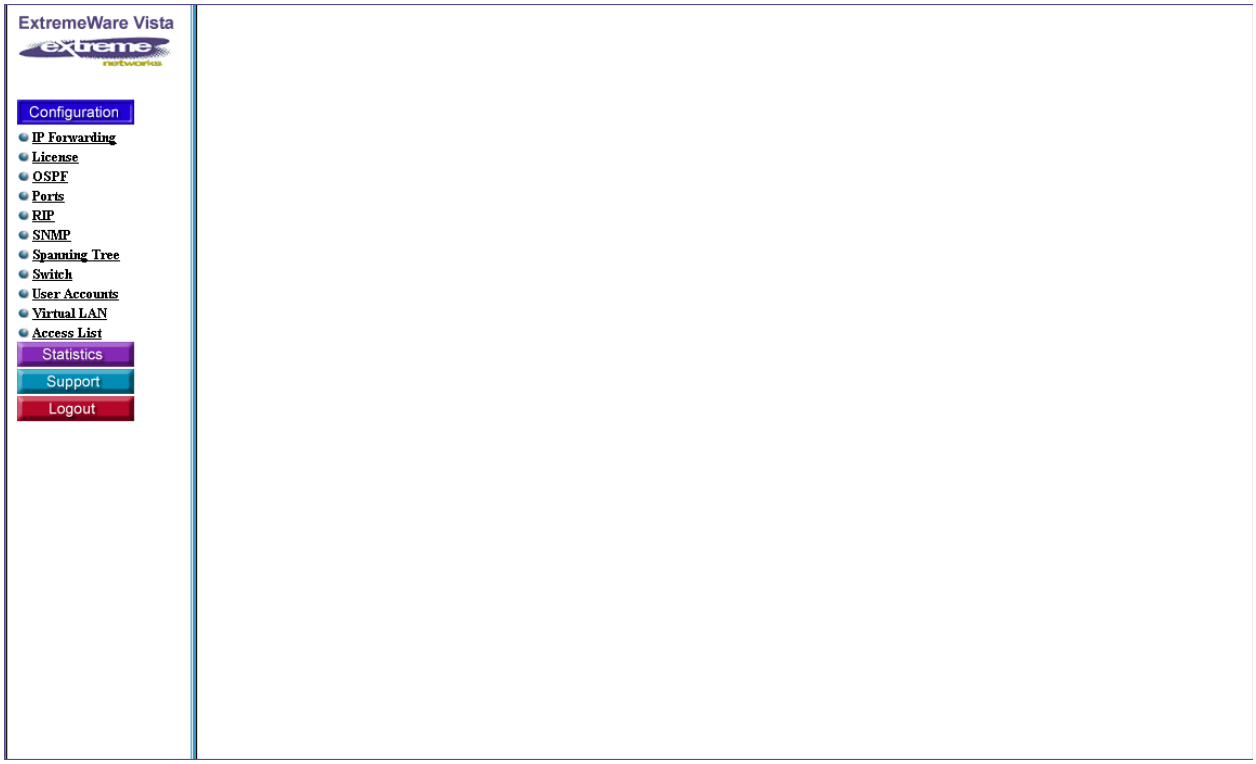

IP Forwarding

From this window, you can enable or disable the IP unicast forwarding across VLANs. For an example of this window, see Figure 53. In the top of the window is a table that shows each existing IP interface configuration. The configuration box that follows allows you to use the pull-down menu to enable or disable forwarding on those existing VLANs. Before submitting a change, users must select the appropriate value for all fields.

The configuration box has the following selectable fields:

VLAN name

Unicast Forwarding—Either enable or disable

Broadcast Forwarding—Either enable or disable

Multicast Forwarding—Enable, disable, or don't change

For more information on forwarding of IP packets, see:

- Configuring IP Unicast Routing on page 220
- Subnet-Directed Broadcast Forwarding on page 218
- IP Multicast Routing Overview on page 245

Figure 53: IP Interface Configuration

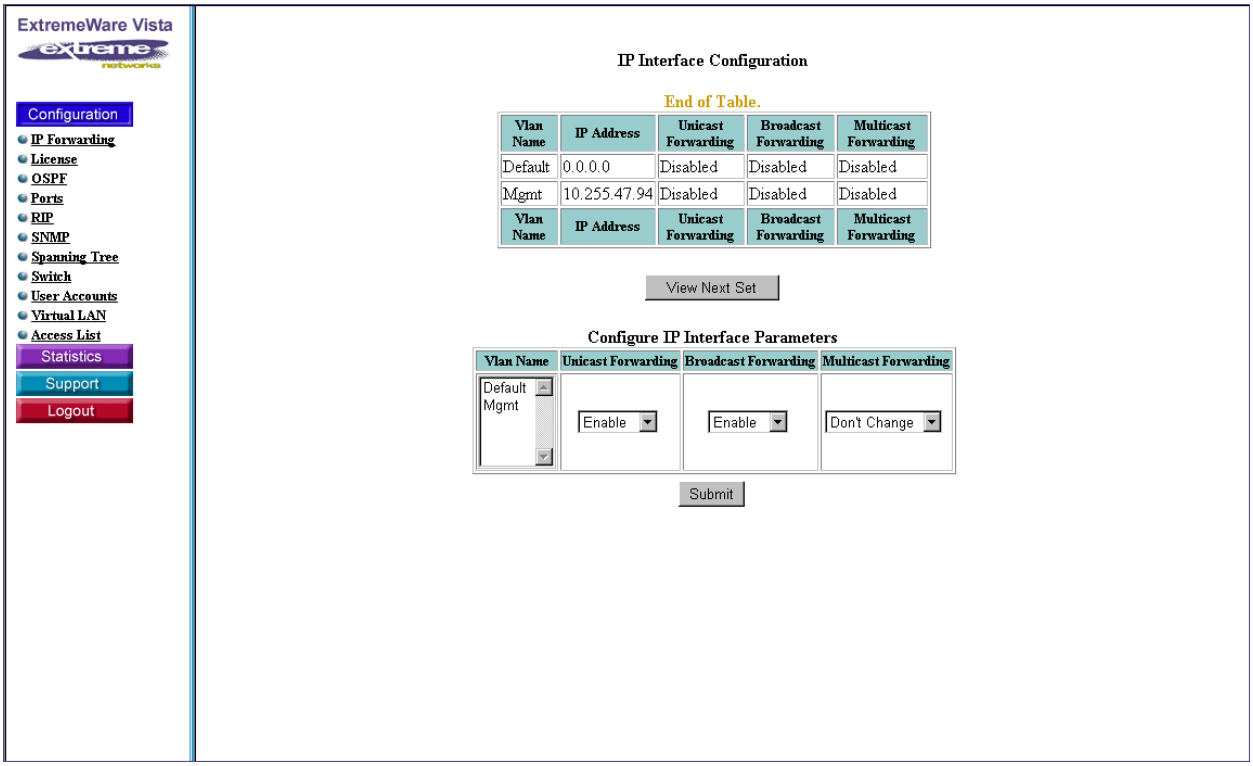

License

The License window allows you to enable the Advanced Edge license by submitting a valid license key purchased from Extreme Networks. See Figure 54 for an example of this window. For more information on levels of licensing, see “Software Licensing” on page 30.

Figure 54: License Window

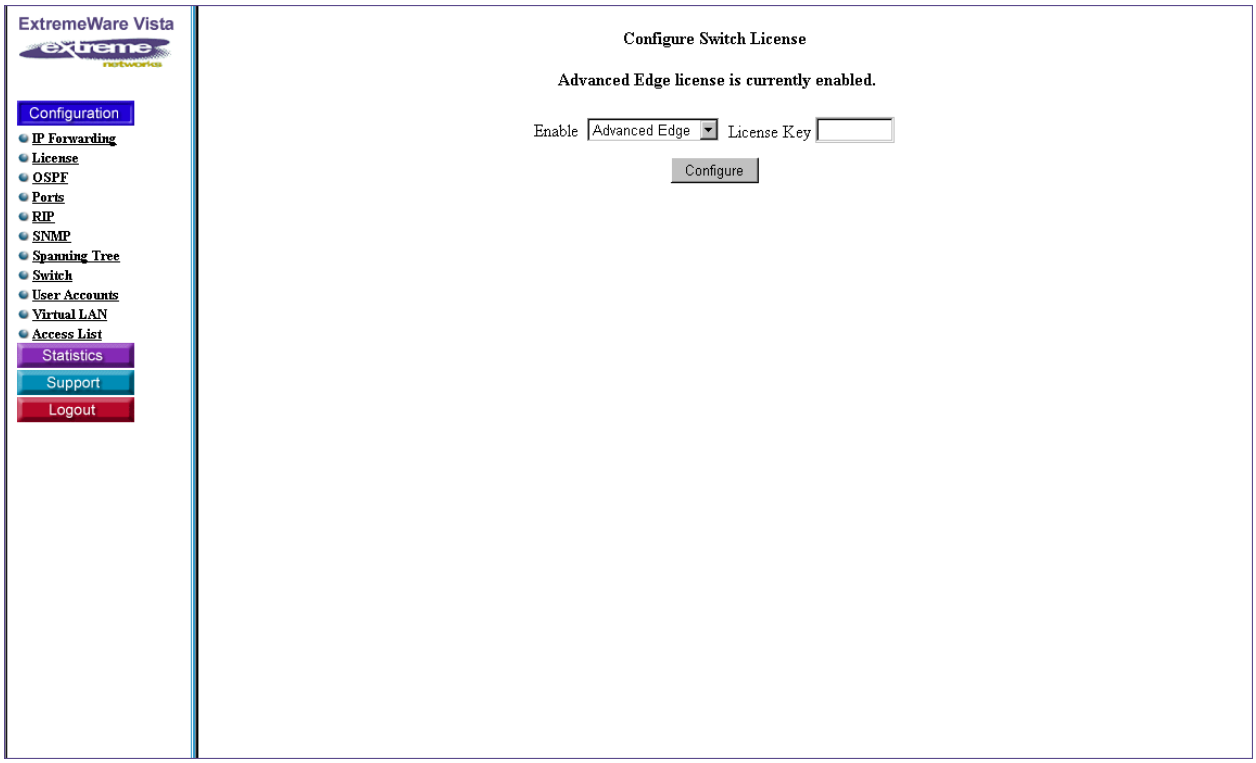

OSPF

The OSPF configuration window allows you to perform a wide-range of OSPF configuration tasks. The window is divided into six functional areas:

1. Configure global OSPF parameters including enabling or disabling of the exporting of RIP, static, and direct (interface) routes to OSPF
2. Create or delete an OSPF area
3. Configure a range of IP addresses in an OSPF area
4. Configure an OSPF area
5. Configure an IP interface for OSPF
6. Configure OSPF authentication

Configure Global OSPF Parameters

Use the global parameters to set up OSPF throughout the switch. See the top portion of Figure 55 for an example of the global parameters window.

NOTE

Before you can make global changes to OSPF, you must first disable OSPF Export Static and OSPF Export RIP.

From this portion of the window, you can:

- Enable or disable the exporting of RIP, static, and direct (interface) routes to OSPF. Be sure you disable exporting of static and RIP before setting other global OSPF parameters.
- Enable or disable the exporting of static, direct, and OSPF-learned routes into a RIP domain.
- Set the route type as external type 1 or external type 2.
- Set the cost metric for all RIP-learned, static, and direct routes injected into OSPF. If the cost metric is set to 0, the cost is inserted from the route.
- Set a tag value for use by special routing applications. Use 0 if you do not have specific requirements for using a tag. The tag value in this instance has no relationship with 802.1Q VLAN tagging.
- Set the OSPF router ID to a user-specified value or to automatic.
- Enable or disable OSPF.

Figure 55: Global OSPF Parameters and Creating or Deleting an Area

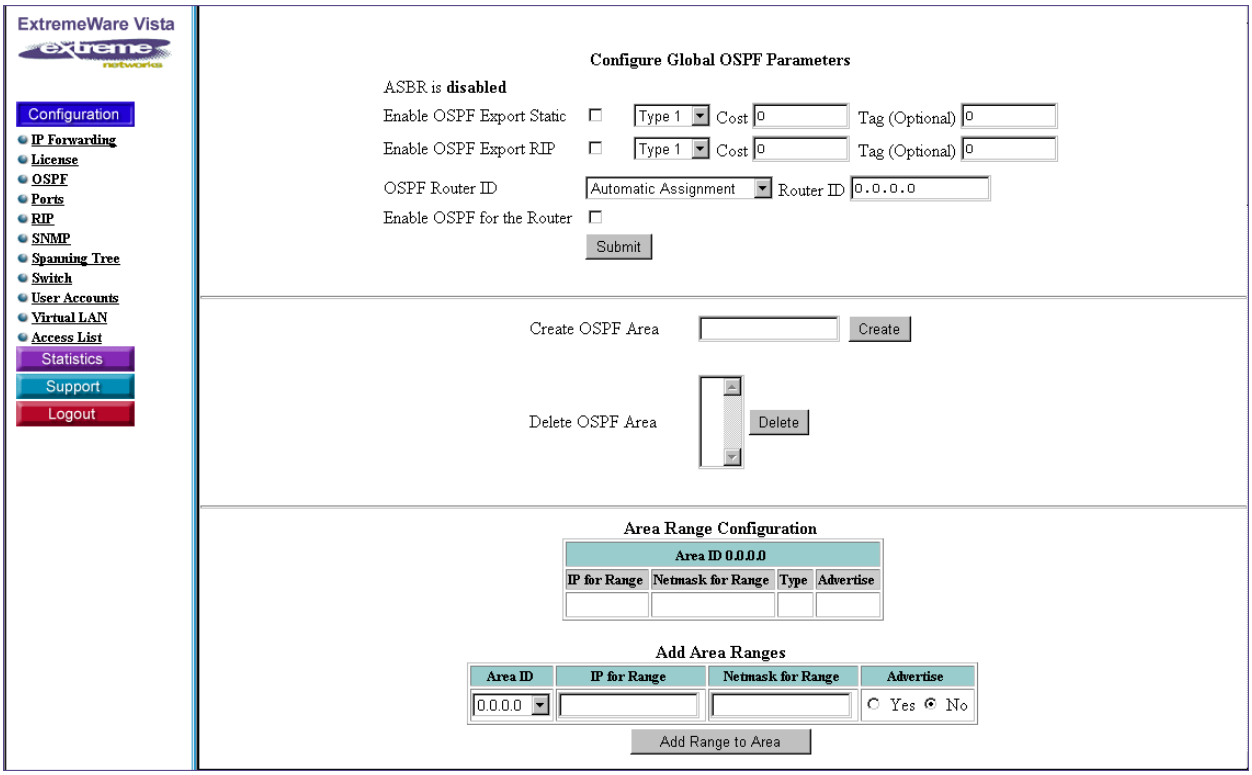

For further details:

- On router IDs, see “Configuring OSPF” on page 238.
- On exporting RIP or OSPF, external types, costs and tags, see “Route Re-Distribution” on page 236.

Create or Delete an OSPF Area

Below the global OSPF parameters is a section dedicated to creating or deleting OSPF areas. Before you configure an area, you must create it. Enter an area ID in the same format as an IP address, (for example, 1.2.3.4).

This portion of the window is also shown in Figure 55. For further details see “Backbone Area (Area 0.0.0.0)” on page 232.

Configure an Area Range

This portion of the window allows you to configure a range of IP addresses in an OSPF area. The example in Figure 56 shows that six areas are defined: the backbone (0.0.0.0), and area IDs 1.1.1.1, 2.2.2.2, 3.3.3.3, 4.4.4.4, and 5.5.5.5. The Area Range Configuration box shows non-default values for the areas. The Add Area Ranges allow you to add a range to an area, set a netmask, or to specify advertising. If advertised, the range is exported as a single LSA by the ABR. You can also delete a range of IP addresses in an OSPF area.

Figure 56: Area Range Configuration

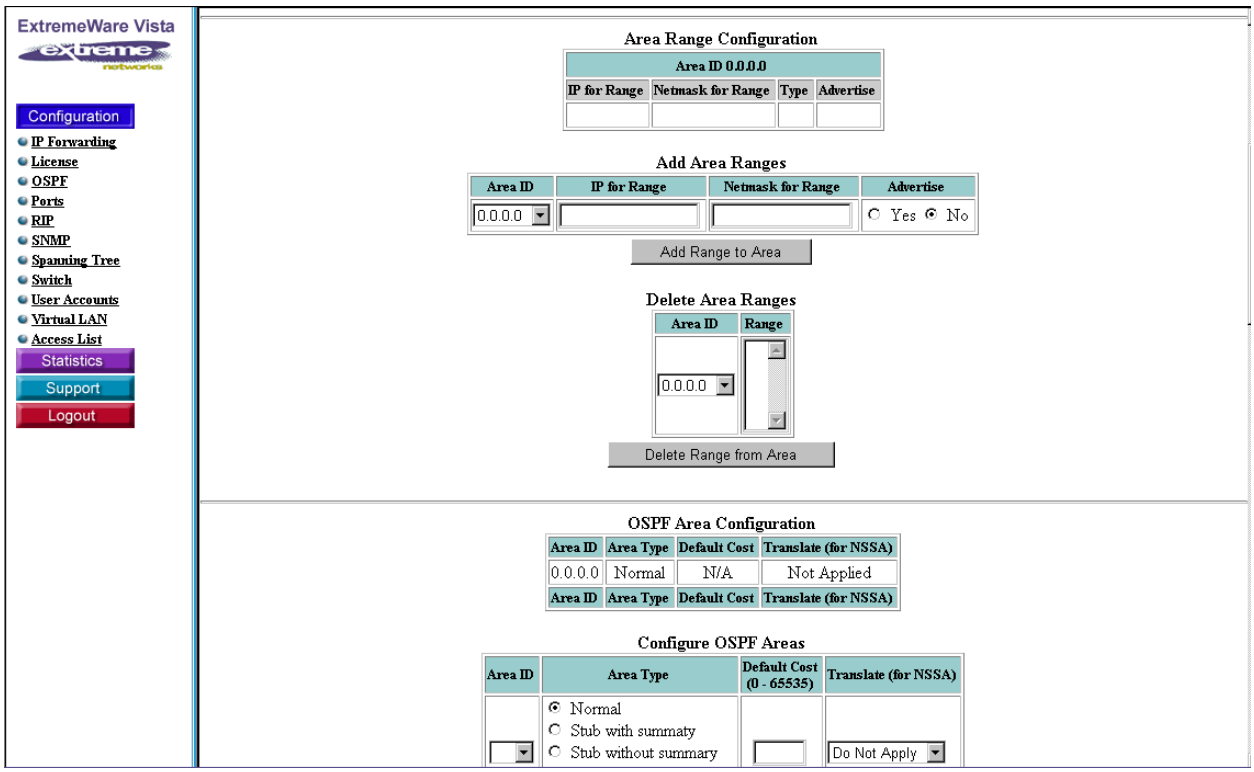

Configure an OSPF Area

Use the scroll bar to locate the next section of the window dedicated to OSPF area configuration, shown in Figure 57. The first table in this section shows each existing configuration. The table that follows allows you to use the pull-down menu to select an area ID. You can also set the area type, the cost, and determine whether to translate for NSSA or not. You may only translate for area type NSSA.

Figure 57: OSPF Area Configuration

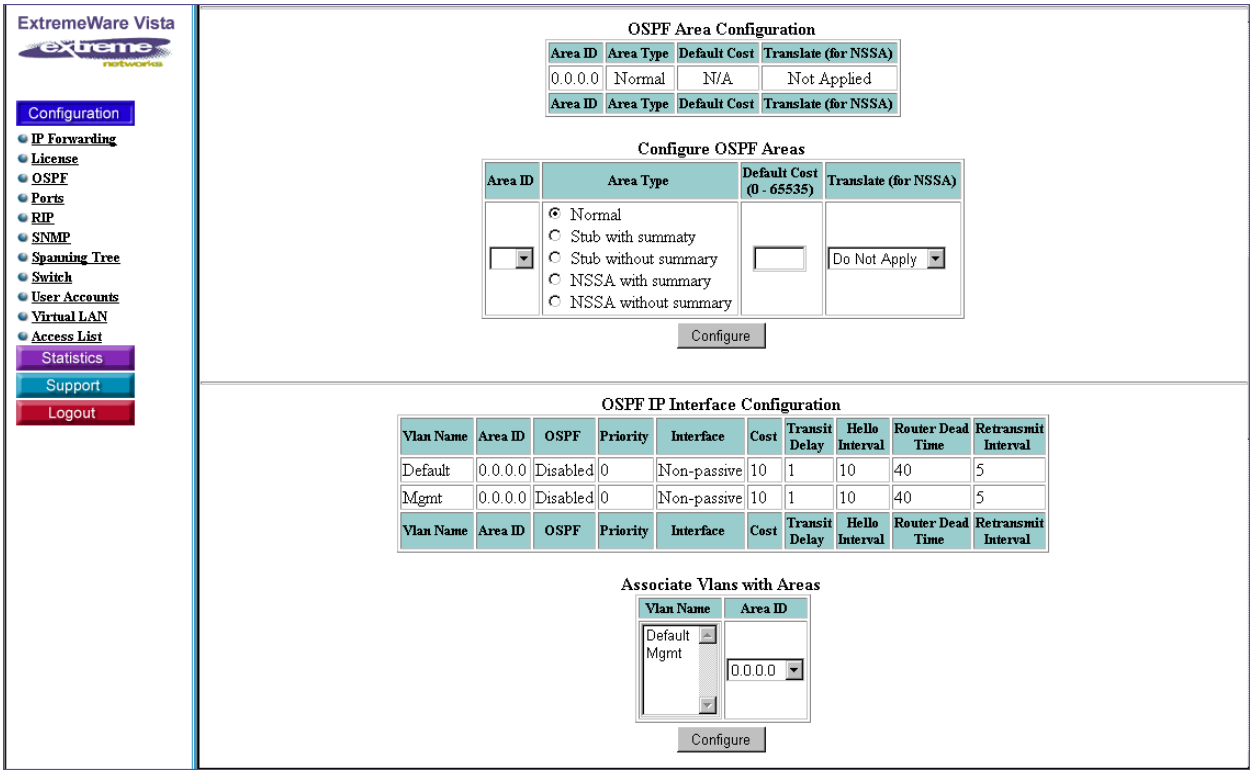

For more information on area types, see “Areas” on page 232.

Configure an IP interface for OSPF

Using this portion of the window, you can:

- Review the existing OSPF IP interface configuration
- Associate a VLAN with an area ID
- Configure OSPF for each VLAN area
- Configure a route filter for non-OSPF routes exported into OSPF
- Configure the timers for one interface in the same OSPF area
- Configure miscellaneous OSPF parameters, such as cost
- Configure virtual links

As shown in Figure 58, the top table lists the existing OSPF IP interface configuration. The table consists of the following fields:

VLAN name

Area ID

OSPF—Either enabled or disabled

Priority—Always set to zero for Summit 400

Interface—Either passive or non-passive

Transit delay—From 1 to 3600 seconds

Hello interval—From 1 to 65535 seconds

Router dead time—From 1 to 2147483647 seconds

Retransmit interval—From 1 to 3600 seconds

The three boxes that follow the table allow you to change the values of the interfaces in that table.

Figure 58: IP Interface Configuration for OSPF

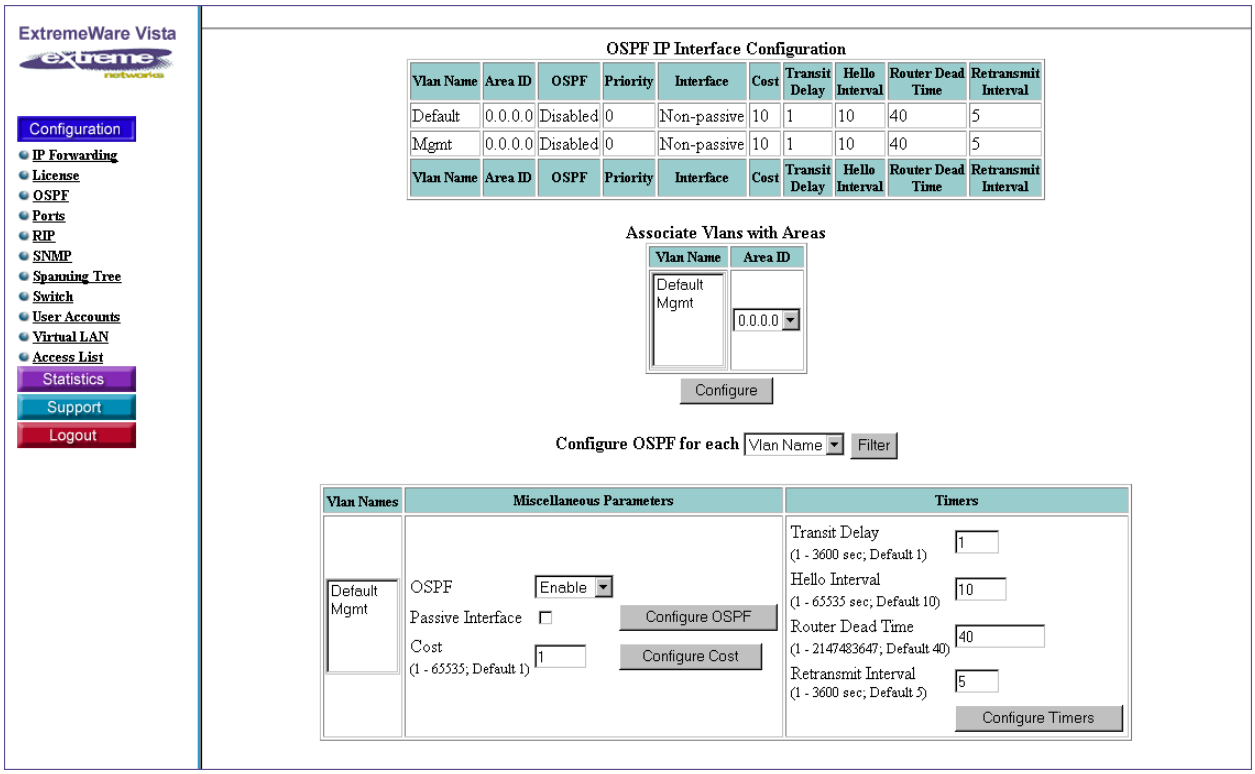

The first box allows you to associate VLANs with areas by selecting a VLAN name and an area ID. The second box allows you to configure OSPF for each VLAN by VLAN name or area ID. The third box, shown in Figure 58 allows you to:

- Select the VLAN by name that is being changed
- Enable or disable OSPF on the interface
- Specify whether the interface is passive or non-passive
- Establish a cost metric
- Set values for timers (transit delay, hello interval, router dead time, and retransmit interval)

Use the next three sets of boxes, shown in Figure 59, to configure virtual links. When non-default values are configured for a router ID or an area ID, the top table displays those values. In the following box you can configure the timers for the virtual link (transit delay, hello interval, router dead time, and retransmit interval).

For further information on virtual links, see “Virtual Links” on page 234.

Figure 59: OSPF Virtual Links

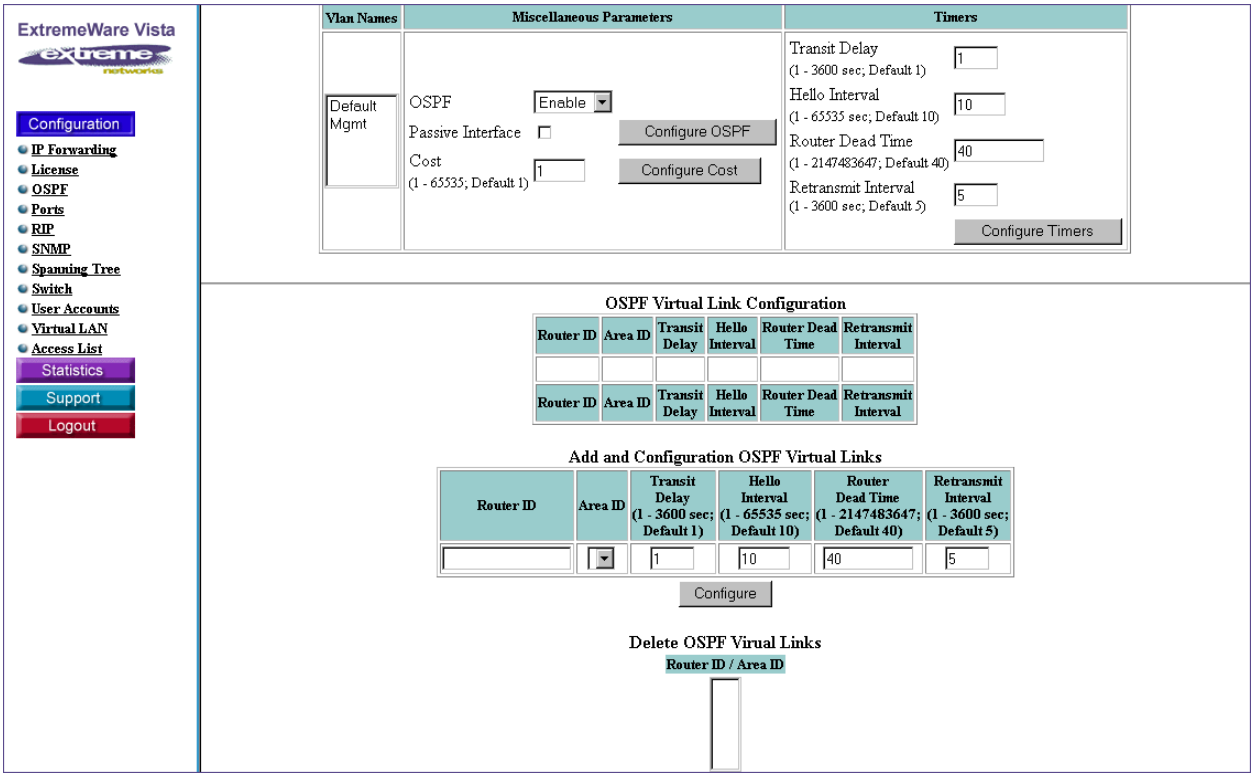

Configure OSPF Authentication

The final section in the OSPF configuration window allows you to configure an interface. This section is shown at the bottom of Figure 60. The table displays the interface and whether an interface type is currently configured. The configuration box allows you to specify a simple authentication password of up to eight characters, or a Message Digest 5 (MD5) key for the interface. If you choose MD5, select a numerical ID between 0 and 255, then select a key value between the range of 0 to 65,535.

Figure 60: OSPF Authentication

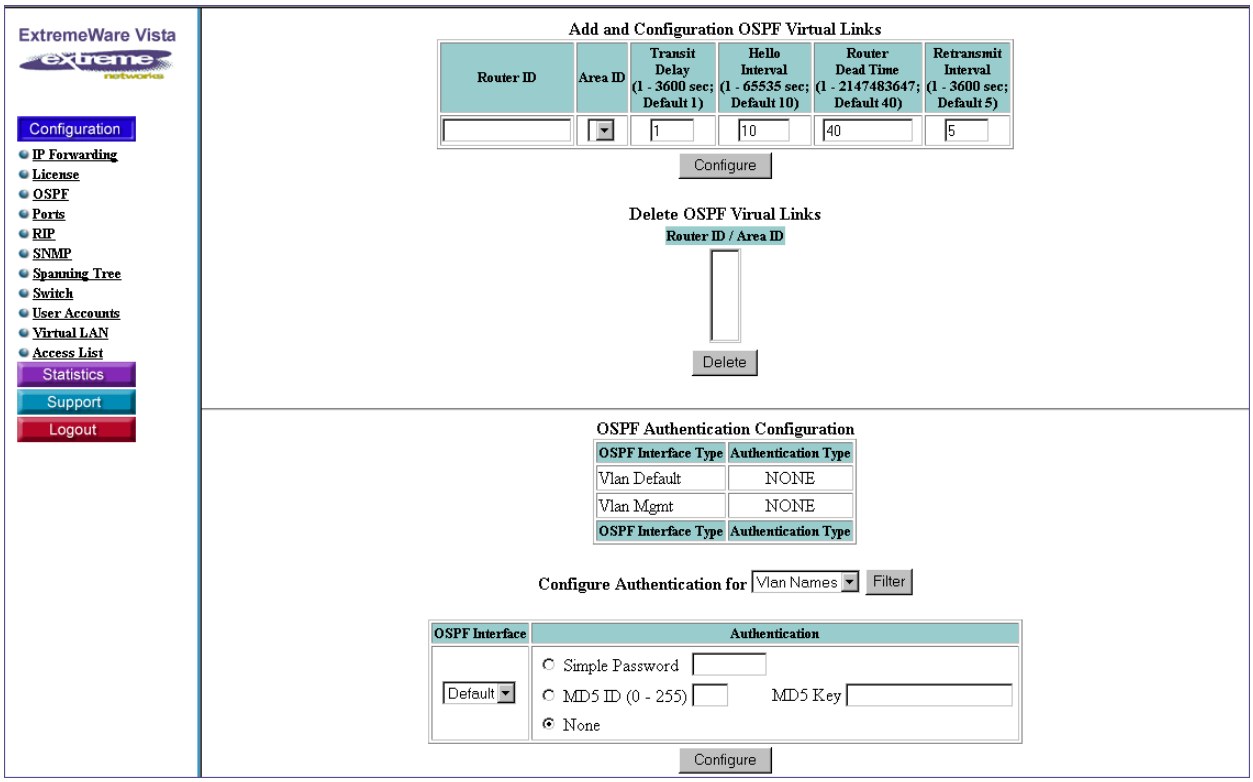

Ports

Port configuration provides a convenient way to see all the pertinent information about a port in one place.

Figure 61 shows the following fields in the port configuration window:

Ports—The port number, 1 to 48

State—The port state, either enabled or disabled

Link—The link status, either active or ready

Autonegotiation—Indicates whether to autonegotiate the port speed and the duplex mode. Autonegotiation is either enabled or disabled.

Configuration Speed—The setting for port speed, either autonegotiated (auto), 10, 100, or 1000

Actual Speed—The speed of the link, either 10, 100, or 1000

Configuration Duplex—The duplex mode, either autonegotiation (auto), half, or full

Actual Duplex—The duplex setting, either half or full

Primary Media—The primary wiring media, either unshielded twisted-pair (UTP) or fiber (SX, LX, or ZX)

Redundant Media—The backup wiring media, always unshielded twisted-pair (UTP)

QoS Profile—A QoS profile in the format of QP*n*, where *n* is from 1 to 8

Figure 61: Port Configuration Window

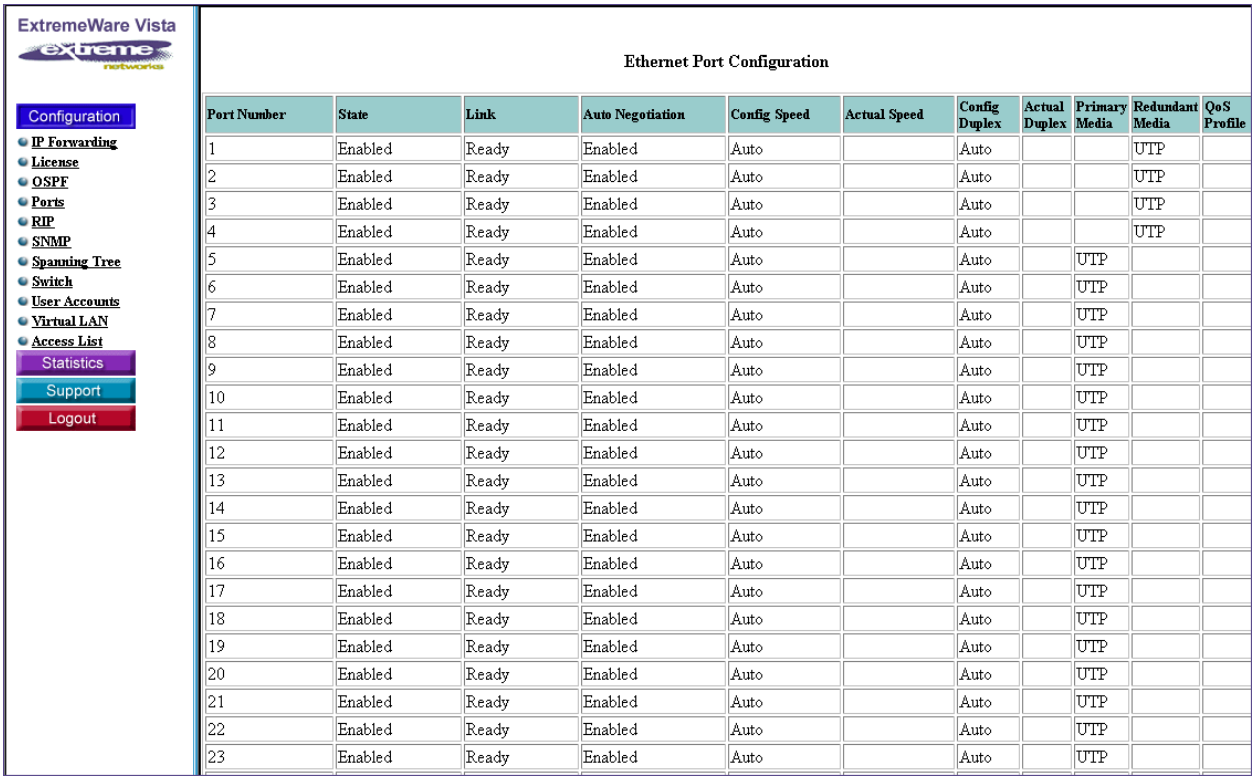

Ethernet Port Configuration

Port Number	State	Link	Auto Negotiation	Config Speed	Actual Speed	Config Duplex	Actual Duplex	Primary Media	Redundant Media	QoS Profile
1	Enabled	Ready	Enabled	Auto		Auto			UTP	
2	Enabled	Ready	Enabled	Auto		Auto			UTP	
3	Enabled	Ready	Enabled	Auto		Auto			UTP	
4	Enabled	Ready	Enabled	Auto		Auto			UTP	
5	Enabled	Ready	Enabled	Auto		Auto		UTP		
6	Enabled	Ready	Enabled	Auto		Auto		UTP		
7	Enabled	Ready	Enabled	Auto		Auto		UTP		
8	Enabled	Ready	Enabled	Auto		Auto		UTP		
9	Enabled	Ready	Enabled	Auto		Auto		UTP		
10	Enabled	Ready	Enabled	Auto		Auto		UTP		
11	Enabled	Ready	Enabled	Auto		Auto		UTP		
12	Enabled	Ready	Enabled	Auto		Auto		UTP		
13	Enabled	Ready	Enabled	Auto		Auto		UTP		
14	Enabled	Ready	Enabled	Auto		Auto		UTP		
15	Enabled	Ready	Enabled	Auto		Auto		UTP		
16	Enabled	Ready	Enabled	Auto		Auto		UTP		
17	Enabled	Ready	Enabled	Auto		Auto		UTP		
18	Enabled	Ready	Enabled	Auto		Auto		UTP		
19	Enabled	Ready	Enabled	Auto		Auto		UTP		
20	Enabled	Ready	Enabled	Auto		Auto		UTP		
21	Enabled	Ready	Enabled	Auto		Auto		UTP		
22	Enabled	Ready	Enabled	Auto		Auto		UTP		
23	Enabled	Ready	Enabled	Auto		Auto		UTP		

Below the Port Configuration table is the box for configuring port parameters. When configuring ports, you must select appropriate values for all parameters before submitting the change. The selectable fields are:

Port Number—Port numbers 1 to 48, or from 1 to 50 if you have the optional XEN card installed.

State—The port state, either enabled or disabled

Restart—Select yes to restart the port

Autonegotiation—The autonegotiation of the port speed and the duplex setting, either enabled or disabled

Speed—The setting for port speed, either 10, 100, or 1000

Duplex—The autonegotiation setting for the duplex setting, either half or full

QoS Profile—A QoS profile in the format of QP*n*, where *n* is from 1 to 8

Figure 62: Configure Port Parameters

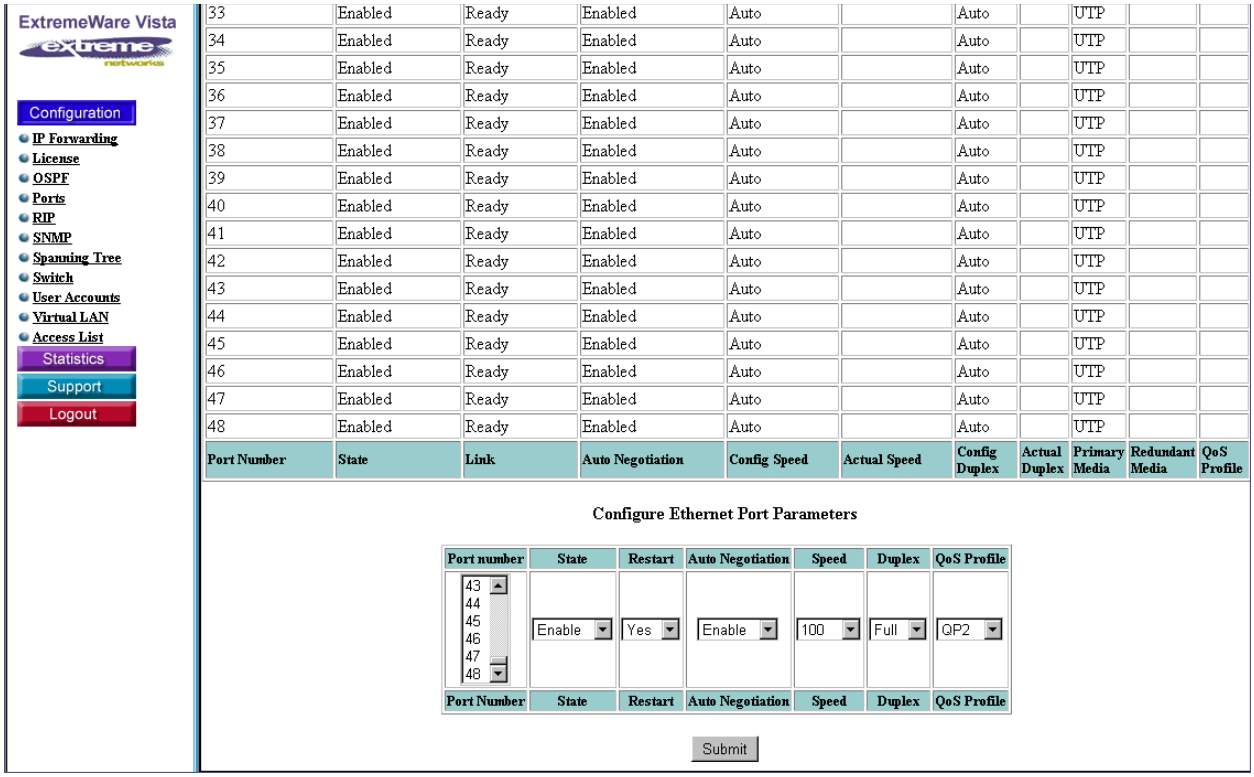

RIP

The RIP configuration window allows you to configure global RIP parameters or RIP for an IP interface.

Configure Global RIP Parameters

Use the global parameters to set up RIP for the switch. See the top portion of Figure 63 for an example of the global parameters window. From this portion of the window, you can make multiple changes with a single update:

- Enable or disable RIP for the switch.
- Enable or disable aggregation.
- Enable or disable redistribution of OSPF static routes through RIP.
- Enable or disable split horizon algorithm for RIP.
- Enable or disable poison reverse algorithm.
- Enable or disable trigger update mechanism.
- Change the periodic RIP update timer.
  - Minimum setting = 10 seconds
  - Maximum setting = Less than the RIP route timeout
  - Default setting = 30 seconds
- Change the route timeout. The default setting is 180 seconds.
- Change the RIP garbage time. The timer granularity is 10 seconds. The default setting is 120 seconds.

Use the **Unconfigure** button to reset the global RIP parameters to the default values. Use the **Submit** button to submit the changes to the system.

Figure 63: RIP Global Configuration

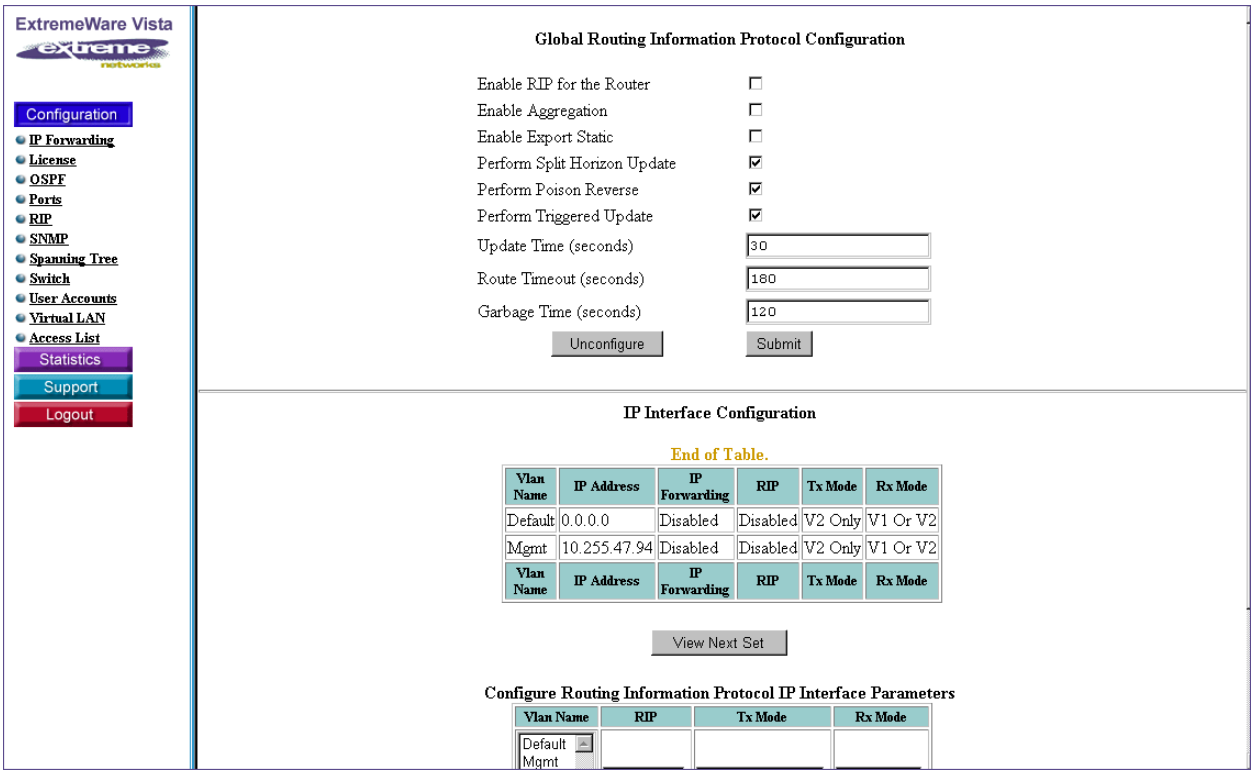

For more information about setting RIP parameters globally, see “Overview of RIP” on page 229.

Configure RIP for an IP interface

Following the global configuration section is for configuring RIP for an individual IP interface. Figure 64 shows an example of this section of the window.

Figure 64: IP Interface Configuration for RIP

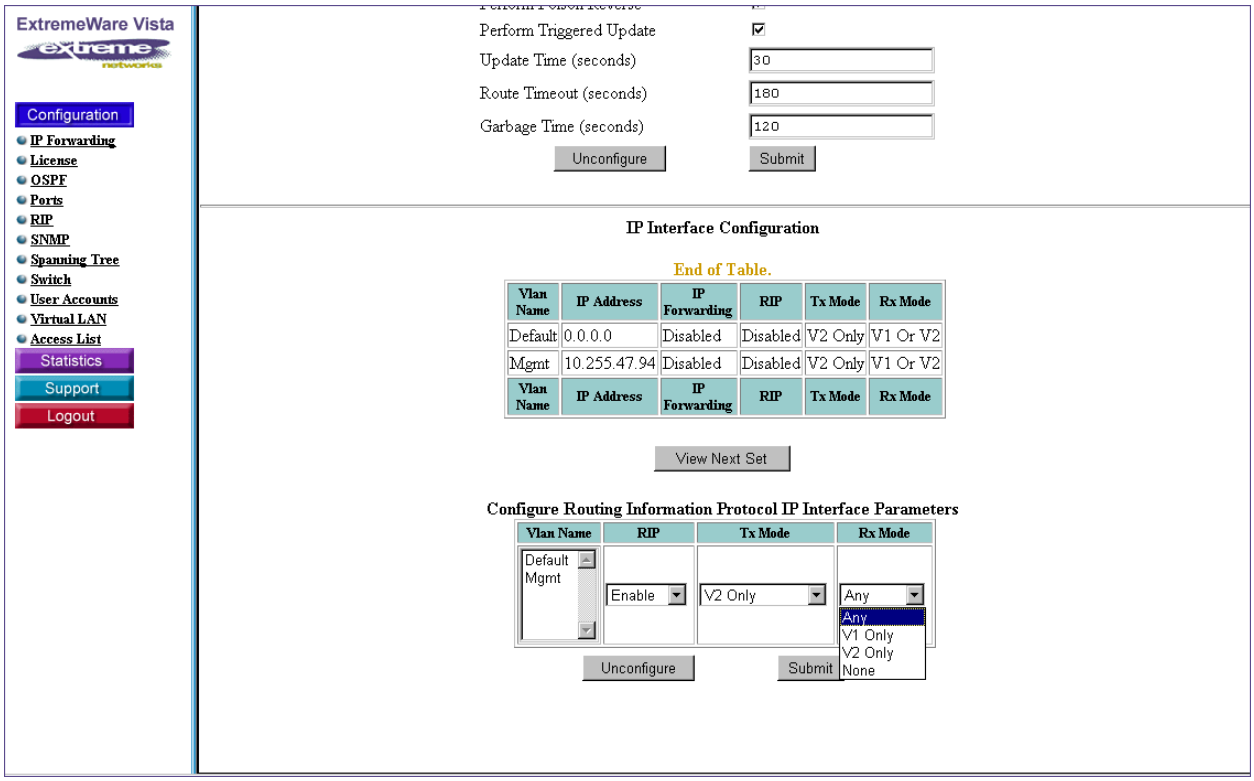

Using this portion of the window, you can:

- Review the existing RIP configuration for an IP interface.

  Each VLAN shows:

  — The VLAN name

  — The IP address

  — Whether IP forwarding is enabled or disabled

  — Whether RIP is enabled or disabled

  — The RIP version used in receive mode (Rx)

  — The RIP version used in transmission mode (Tx)
- Enable or disable RIP on a VLAN
- Configure RIP on a VLAN
- Set the Tx mode values for the selected VLANs. The pull-down menu allows you to specify the following:

  None—Do not transmit any packets on this interface.

  V1 Only—Transmit RIP v1 format packets to the broadcast address.

  V1 Compatible—Transmit RIP v2 format packets to the broadcast address.

  V2 Only—Transmit RIP v2 format packets to the RIP multicast address.

  If no VLAN is specified, the setting is applied to all VLANs. The default setting is V2 Only.

- Set the Rx mode values for the selected VLANs. The pull-down menu allows you to specify the following:

  None—Do not receive packets on this interface.

  Any—Receive packets on this interface in any mode.

  V1 Only—Receive RIP v1 format packets to the broadcast address.

  V2 Only—Receive RIP v2 format packets to the RIP multicast address.

  If no VLAN is specified, the setting is applied to all VLANs. The default setting is V2 Only.
- Use the **Unconfigure** button to reset the RIP configuration for the VLAN to the default values.
- Use the **Submit** button to submit the changes to the system.

SNMP

The SNMP window is divided into two sections. The top section allows you to enter system group information and authentication information for the community strings. The bottom section allows you to set the configuration associated with SNMP traps.

System Group Configuration

As shown in Figure 65, this portion of the SNMP window allows you to set:

Contact —A text field that enables you to enter the contact information of the person responsible for managing the switch.

Name—The system name is the name that you have assigned to this switch. The default name is the model name of the switch (for example, Summit 400-48t switch).

Location —The location of this switch.

Figure 65: System Contact and Community Authentication Information

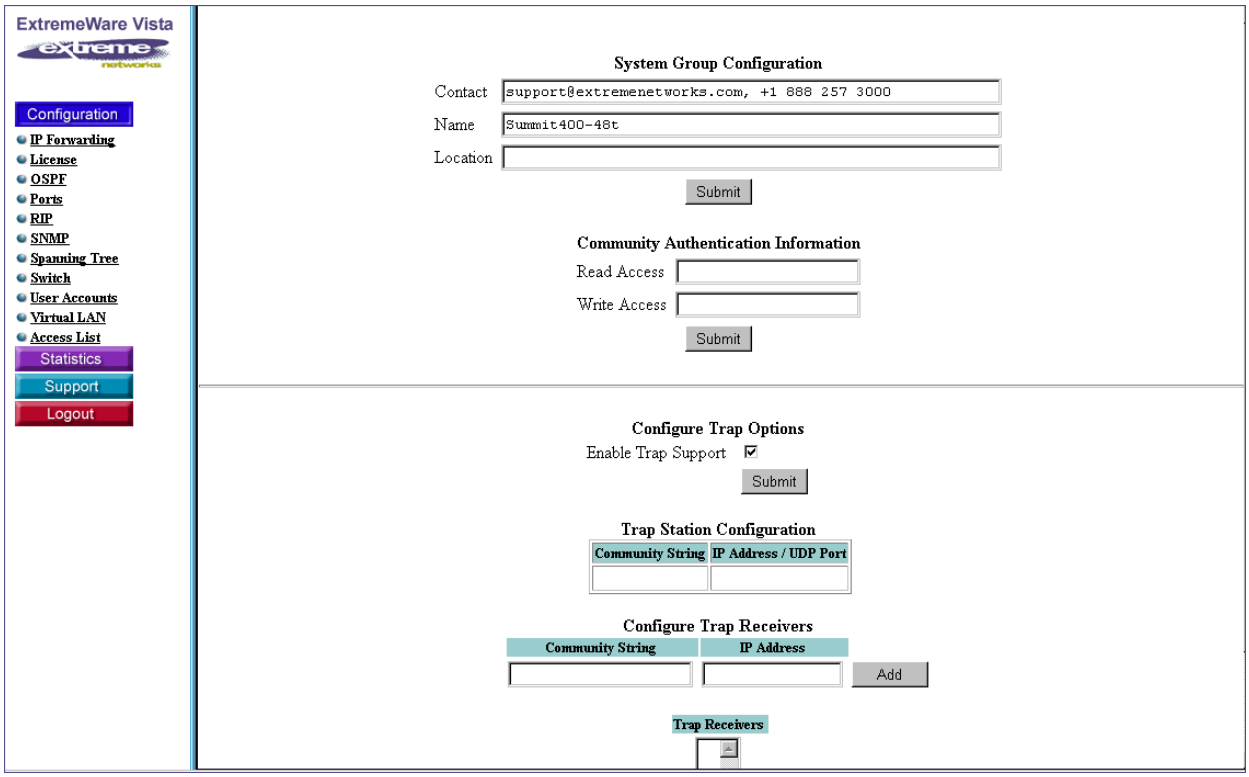

The Community Authentication Information fields specify community strings, which allow a simple method of authentication between the switch and the remote Network Manager. The default read-only community string is `public`. The default read-write community string is `private`. Each community string can have a maximum of 127 characters, and can be enclosed by double quotation marks.

Trap Information

As shown in Figure 66, the lower section of the SNMP window allows you to enable SNMP and configure trap receivers.

To enable SNMP trap support, click the checkbox and submit the request.

If authorized trap receivers are currently configured on the network, the Trap Station Configuration table lists the community string and IP address or User Datagram Protocol (UDP) port of the trap receivers.

The last two boxes in the section allow you to add a trap receiver or to delete a trap receiver. For further information on SNMP and trap receivers, see “Using SNMP” on page 48.

Figure 66: Configure Trap Options

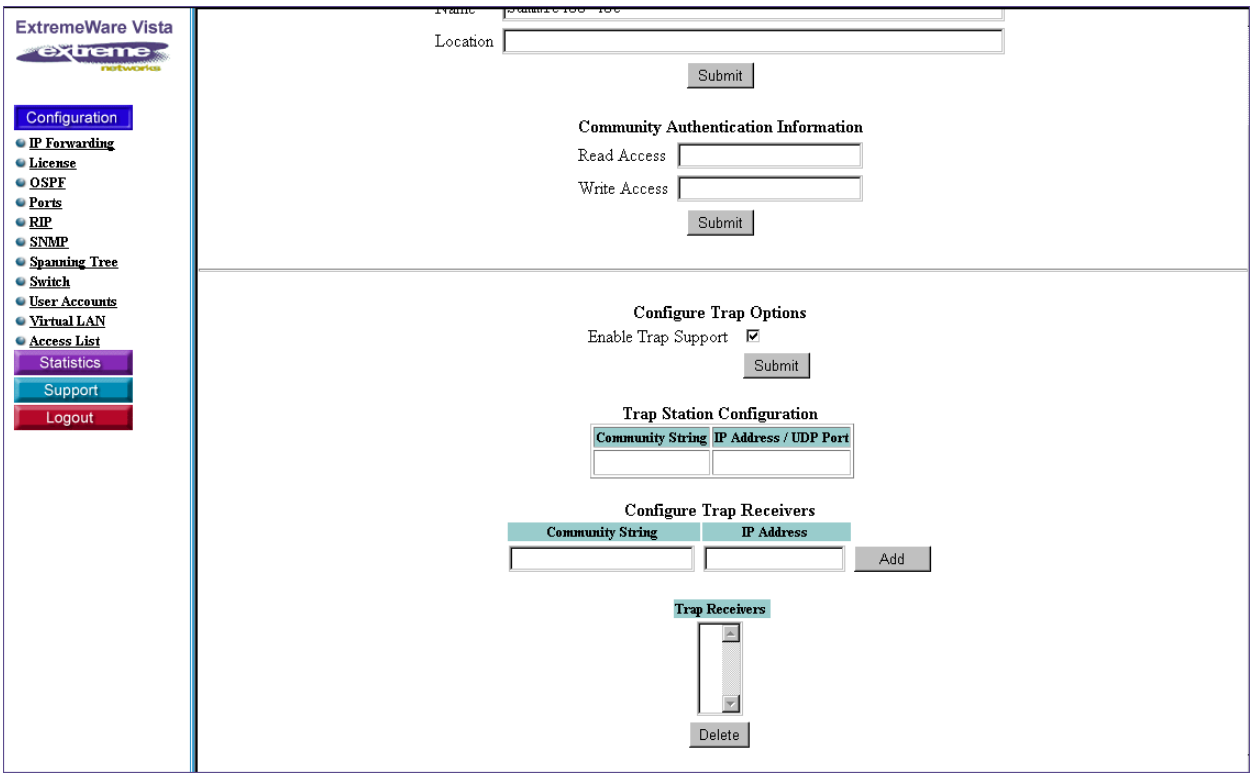

Spanning Tree

From this window, you can configure all aspects of a Spanning Tree Domain (STPD). The window is divided into two sections.

In the top section, you can create or delete a Spanning Tree Domain (STPD) as shown in Figure 67.

Figure 67: Spanning Tree Configuration (1 of 4)

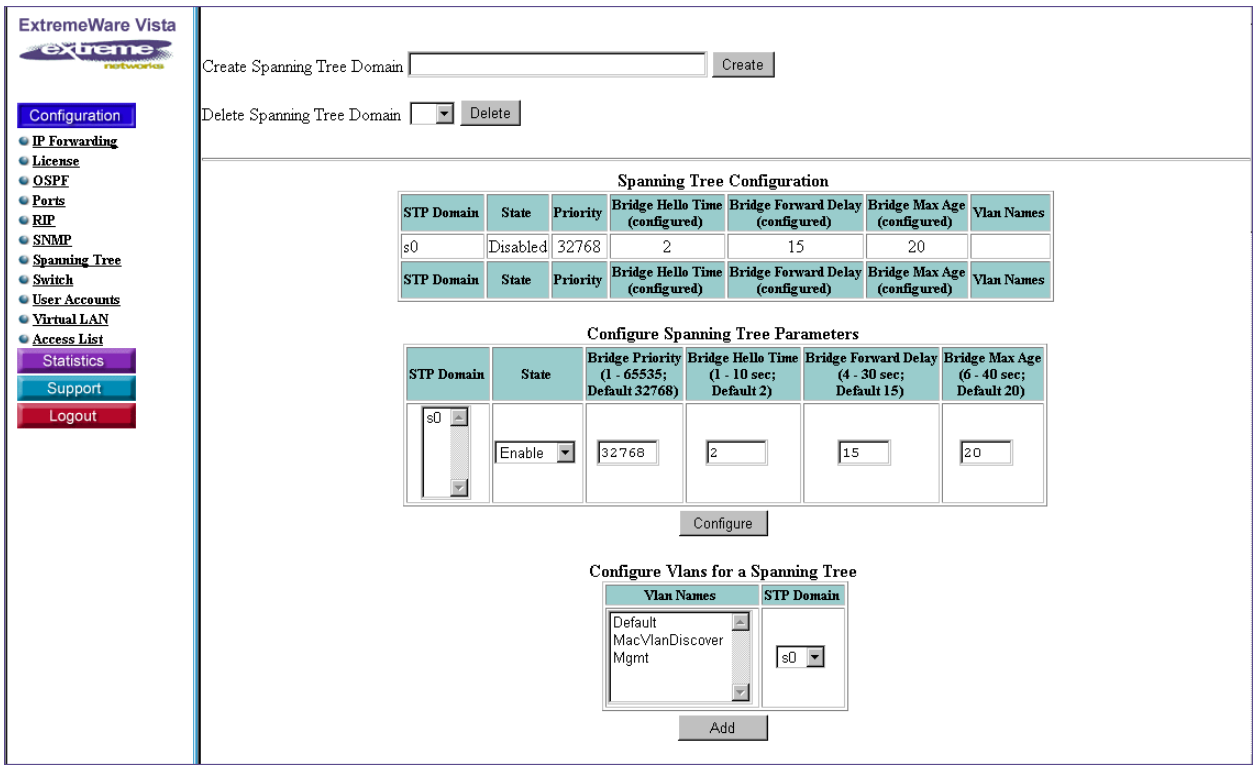

In the bottom section, you can:

- Review all STPD configurations

  Each STPD shows the:

  — STPD name.

  — State of the domain, either enabled or disabled.

  — Priority level of the bridge, a value between 1 and 65535 (default 32768).

  — Hello time interval for the bridge, a value between 1 and 10 seconds (default 2 seconds). The hello time specifies the time delay between the transmission of Bridge Protocol Data Units (BPDUs) from this STPD when it is the Root Bridge.

  — Bridge forward delay, a value between 4 and 30 seconds (default 15 seconds). The bridge forward delay specifies the time that the ports in this STPD spend in the listening and learning states when the switch is the Root Bridge.

  — The maximum age of a BPDU, a value between 6 and 40 seconds (default 20 seconds).

  The STPD configuration table is shown in Figure 67 and Figure 68.

- Create or change parameters on a STPD.

  Select a STPD, change the parameter values as described above, and click **Configure**.

  The Configure Spanning Tree Parameters box is shown in Figure 67 and Figure 68.

- Assign VLANs to a STPD, as shown in Figure 68.
- Unconfigure STPD, as shown in Figure 68.

Figure 68: Spanning Tree Configuration (2 of 4)

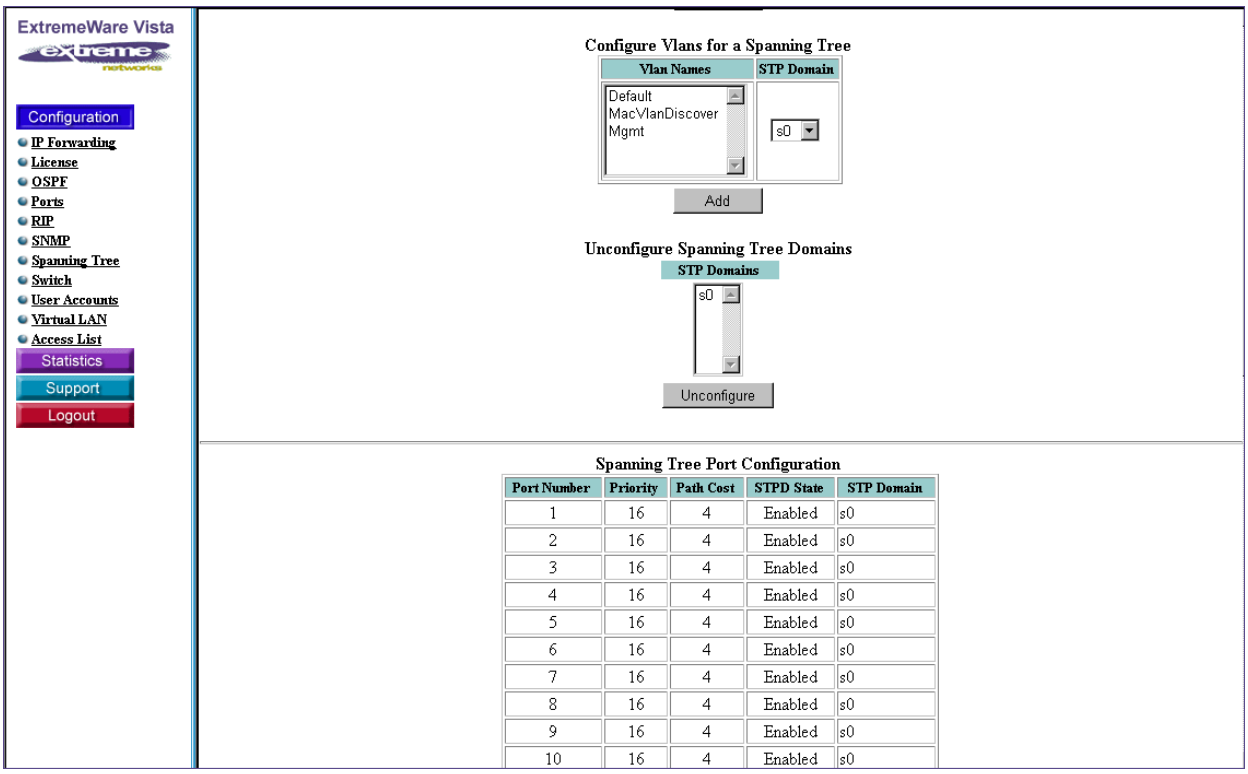

- Review all ports belonging to STPDs.

  A port can belong to only one STPD. If a port is a member of multiple VLANs, then all those VLANs must belong to the same STPD. The Spanning Tree Port Configuration Table contains the following fields:

  Port Number—Port numbers 1 to 48 or from 1 to 50 if you have the optional XEN card installed.

  Priority— The priority of the port indicates the likelihood of the port becoming the root port. The range is 0 through 31, where 0 indicates the lowest priority. The default setting is 16.

  Path Cost—Specifies the path cost of the port in this STPD. The range is 1 through 65,535. The switch automatically assigns a default path cost based on the speed of the port, as follows:

  — For a 10 Mbps port, the default cost is 100.

  — For a 100 Mbps port, the default cost is 19.

  — For a 1000 Mbps port, the default cost is 4.

  — For a 10000 Mbps port, the default cost is 2.

  STPD State—Specifies whether the Spanning Tree Protocol is enabled or disabled on the STPD.

  STP Domain—The name of the STP domain.

  See Figure 69 for an example of the table.

- Configure Spanning Tree ports.

  Add or change the above parameters for STP ports. See Figure 70 for an example of this configuration box.

Figure 69: Spanning Tree Configuration (3 of 4)

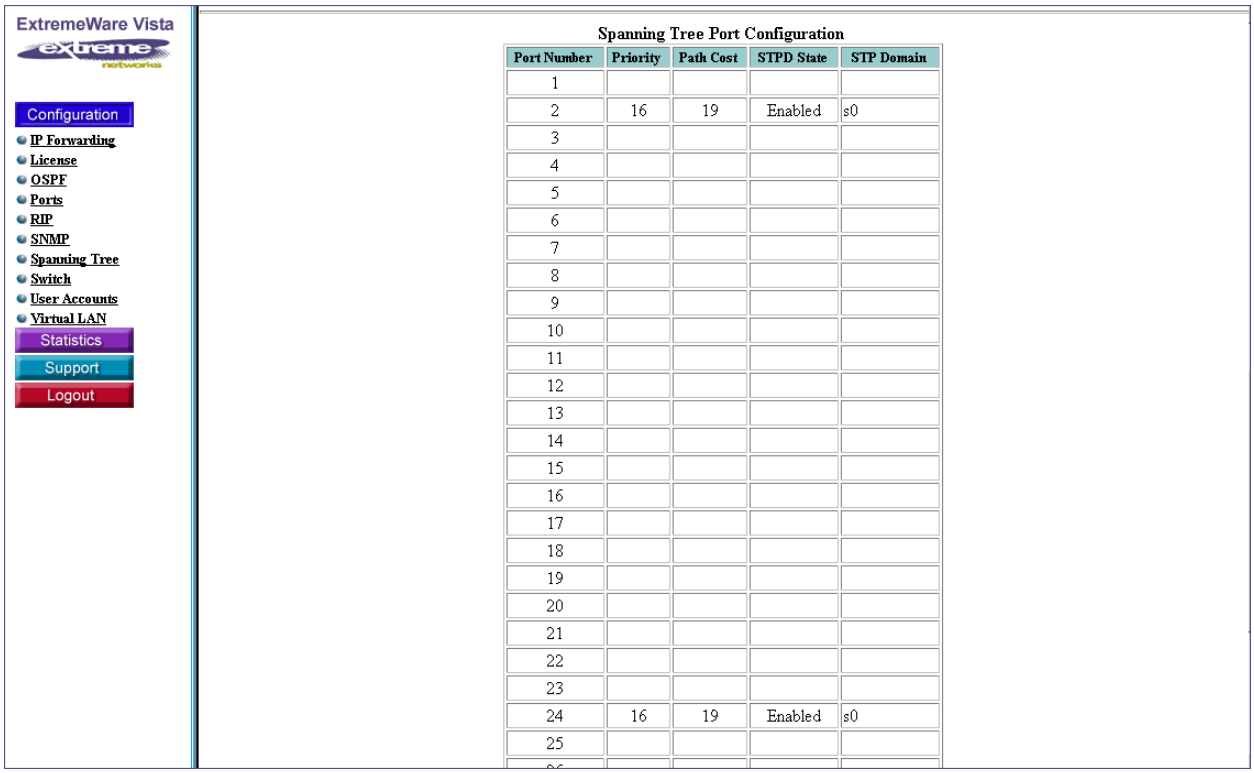

Figure 70: Spanning Tree Configuration (4 of 4)

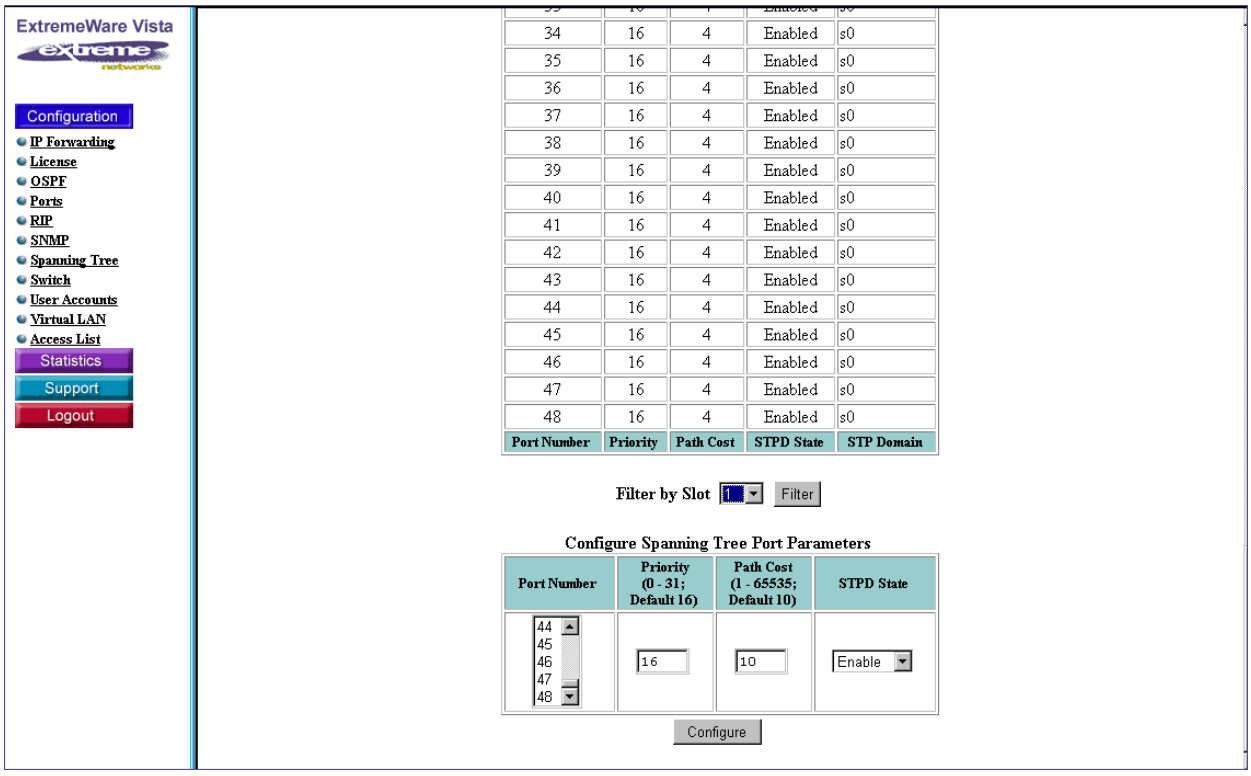

Switch

This window, shown in Figure 71, manages basic switch operation. The four sections are:

- Set date and time
- Enable or disable Telnet remote management and SNMP management
- Select the image and configuration to use

  You can choose a primary or secondary image to use from the pull-down menu.
- Save the configuration

  Settings that are stored in run-time memory are not retained by the switch when the switch is rebooted. To retain the settings, and have them load when you reboot the switch, you must save the configuration to nonvolatile storage.

  The switch can store two different configurations: a primary and a secondary. When you save configuration changes, you can select into which configuration area you want the changes saved. If you do not specify the configuration area, the changes are saved to the configuration area currently in use.
- Reboot the switch

  This stand-alone button causes the Summit 400 to reboot immediately.

Figure 71: Switch Configuration

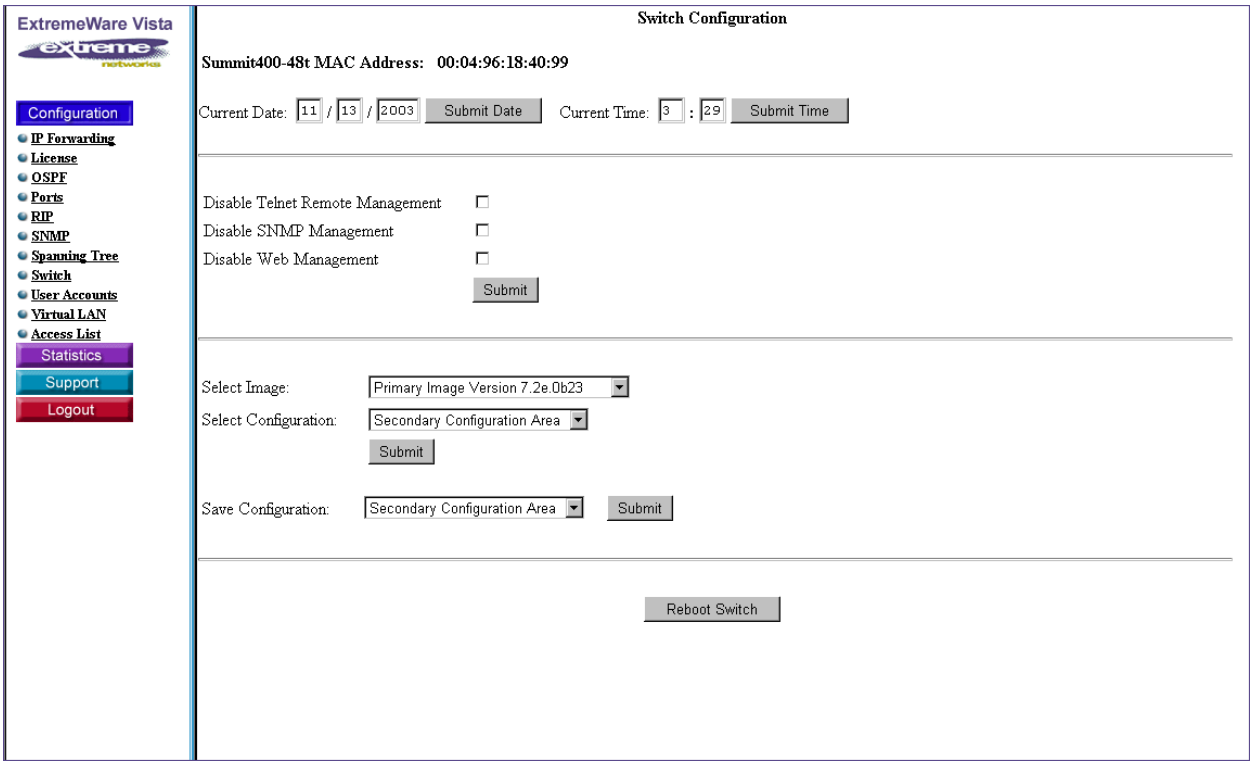

User Accounts

This window allows you to control access to the system. As shown in Figure 72, the top table provides:

- The user's name

- Whether that user has administrator privileges
- The number of times the user has logged into the system since the last reboot
- Whether the user has point-to-point (PPP) user access

You can also manage user accounts through this window. Each account requires a user name and password. Users with administrative access have read-write authority, where normally a user would have read-only access to the system. Only users with read-write authority have permission to change the switch's configuration. There is also a checkbox to delete a user.

For more information on controlling user access, see "Configuring Management Access" on page 70.

Figure 72: Management Access

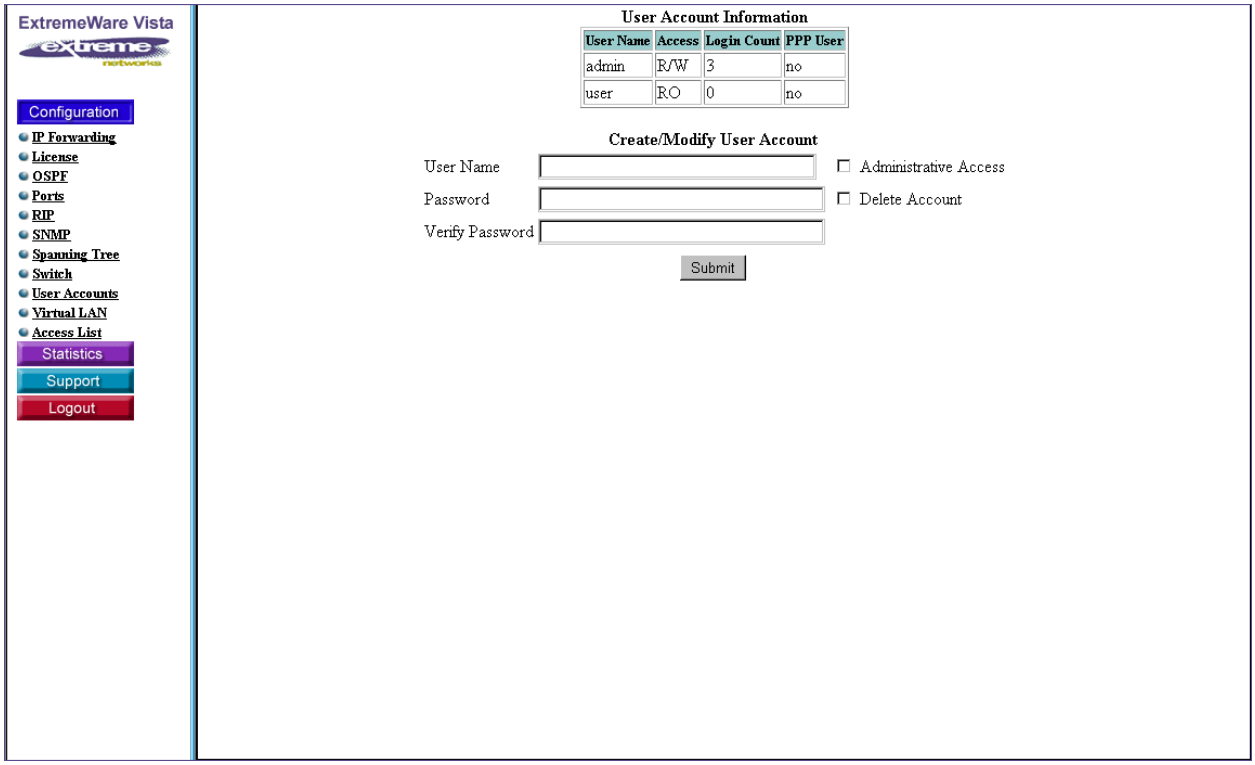

Virtual LAN

This window allows you to perform the most common VLAN administration tasks. It is divided into three sections:

- Creating and deleting a VLAN
- Changing a VLAN name
- Configuring a VLAN

Creating and Deleting a VLAN

The top section of the window allows you to create or delete a VLAN, as shown in Figure 73. When naming a VLAN, be sure to following the naming guidelines described in "VLAN Names" on page 92.

Figure 73: VLAN Administration (1 of 2)

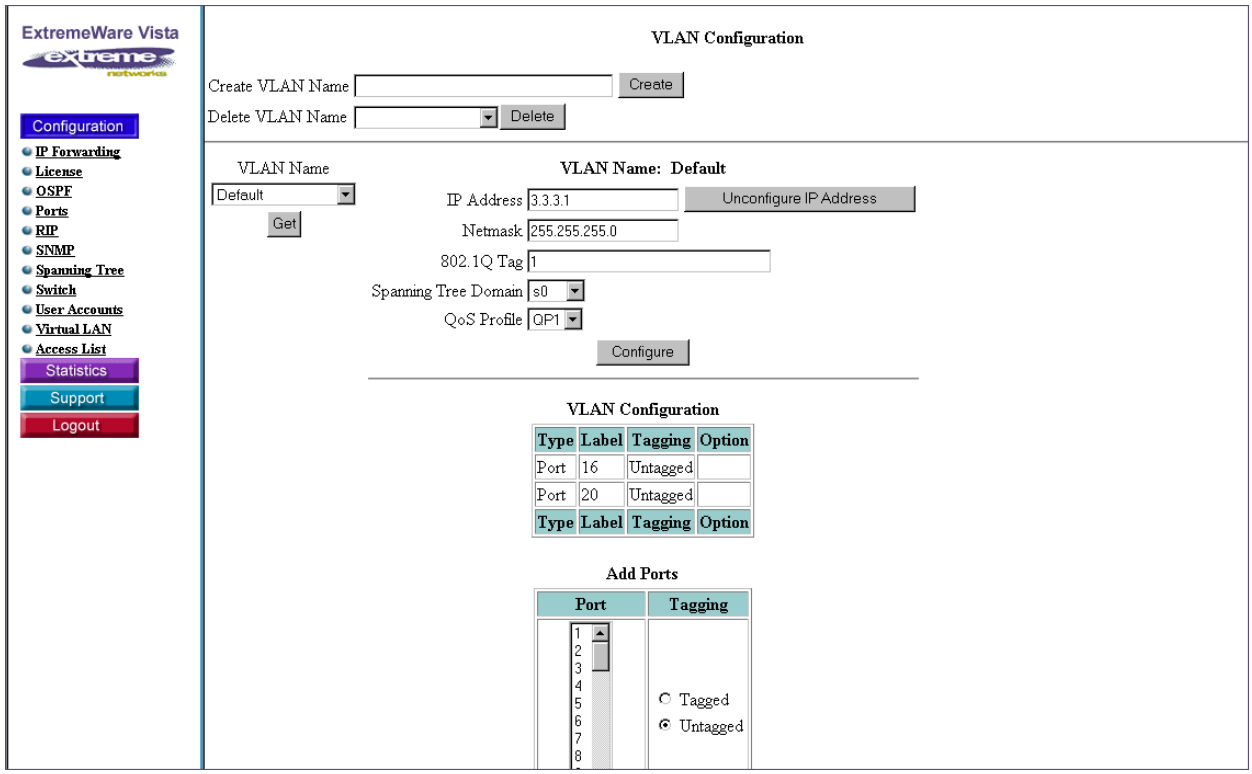

Configuring a VLAN

The second section of the VLAN window allows you to change VLAN parameters. Use the pull-down menu to choose an existing VLAN name and click **Get** to populate the remaining fields. Figure 74 shows an example of the Configure VLAN Information.

Use the following fields to make changes to a VLAN:

IP Address—Either changes the IP address or unconfigures the IP address. The **Unconfigure** button resets the IP address of the VLAN; the **Configure** button allows you to assign a different IP address to the VLAN.

Netmask—Specifies a subnet mask in dotted-quad notation (e.g. 255.255.255.0).

802.1Q Tag—Adds an 802.1Q tag to the VLAN. Acceptable values range from 1 to 4094.

Spanning Tree Domain—Assigns the VLAN to a STPD.

QoS Profile—Assigns a QoS profile to the VLAN.

Figure 74: VLAN Administration (2 of 2)

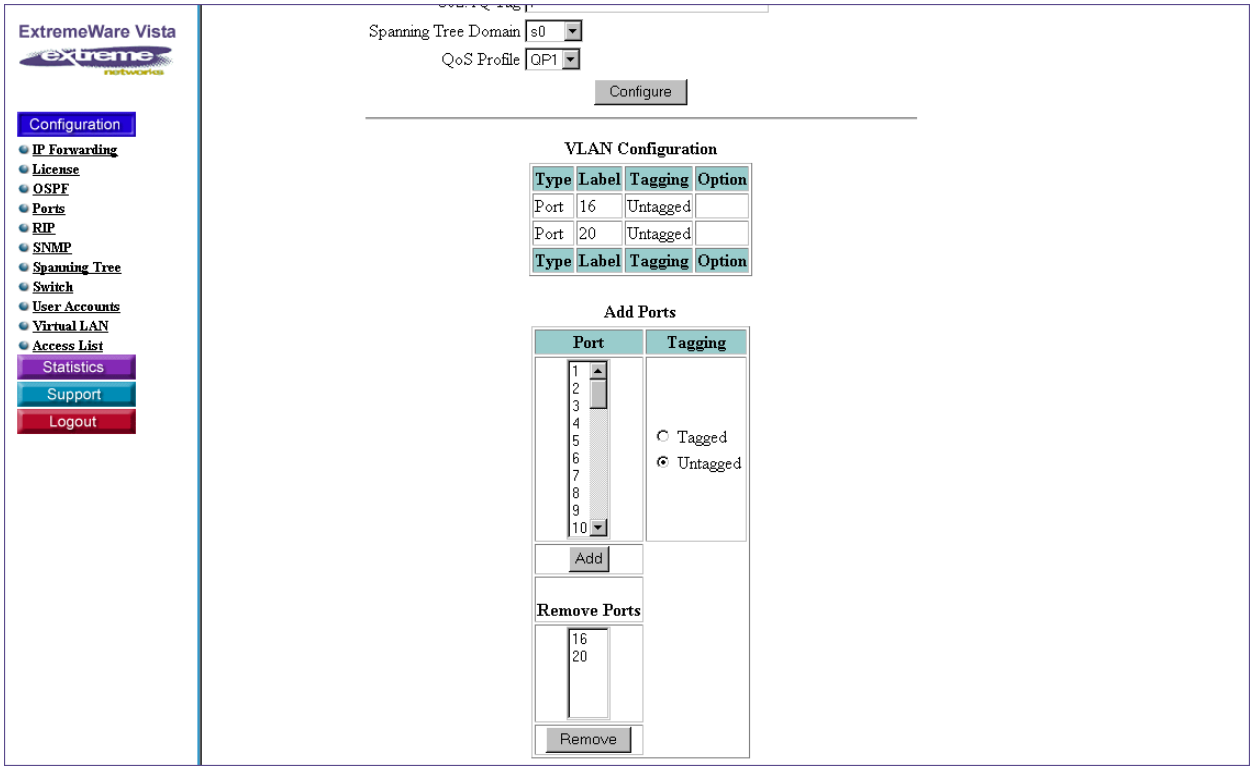

The next section allows you to adds ports to the VLAN.

Adding Ports to a VLAN. You can either add the port as tagged or untagged. If you click **Tagged**, the port is added as a tag-based port. If you click **Untagged**, the port is added as an untagged port.

Figure 74 shows an example of adding ports to a VLAN.

The next box allows you to select a port and click **Remove** to delete the port.

Access List

This window allows you to configure an IP access list, a rate limit, and their associated access masks. IP access lists, also known as access control lists (ACLs) are used to perform packet filtering and forwarding decisions on incoming traffic. Each packet arriving on an ingress port is compared to the access list in sequential order and is either forwarded to a specified QoS profile or dropped.

Each access control list consists of an access mask that selects which fields of each incoming packet to examine, and a list of values to compare with the values found in the packet. Access masks can be shared multiple access control lists, using different lists of values to examine packets.

The top section of the window, as shown in Figure 75, displays information about existing access masks. The following mask features are shown in a table format:

Dest Mac—Ethernet destination MAC address

Src Mac—Ethernet source MAC address

VLAN ID—VLAN identifier (VLANid)

Ether Type—Ethernet type

IP Proto—IP protocol

TOS/Code Point—IP DiffServ code points

Dest IP—Destination IP address

Dest IP Mask—Destination subnet mask

Dest L4 Port—Destination UDP layer 4 port

Src IP—Source IP address

Src IP Mask—Source IP subnet mask

Src L4 Port/ICMP—Source UDP layer 4 port / ICMP

TCP Permit Estb—TCP permit established

Egr Port—Egress port

Ingr Port—Ingress port

Pre—Precedence

Figure 75: Access List Configuration (1 of 3)

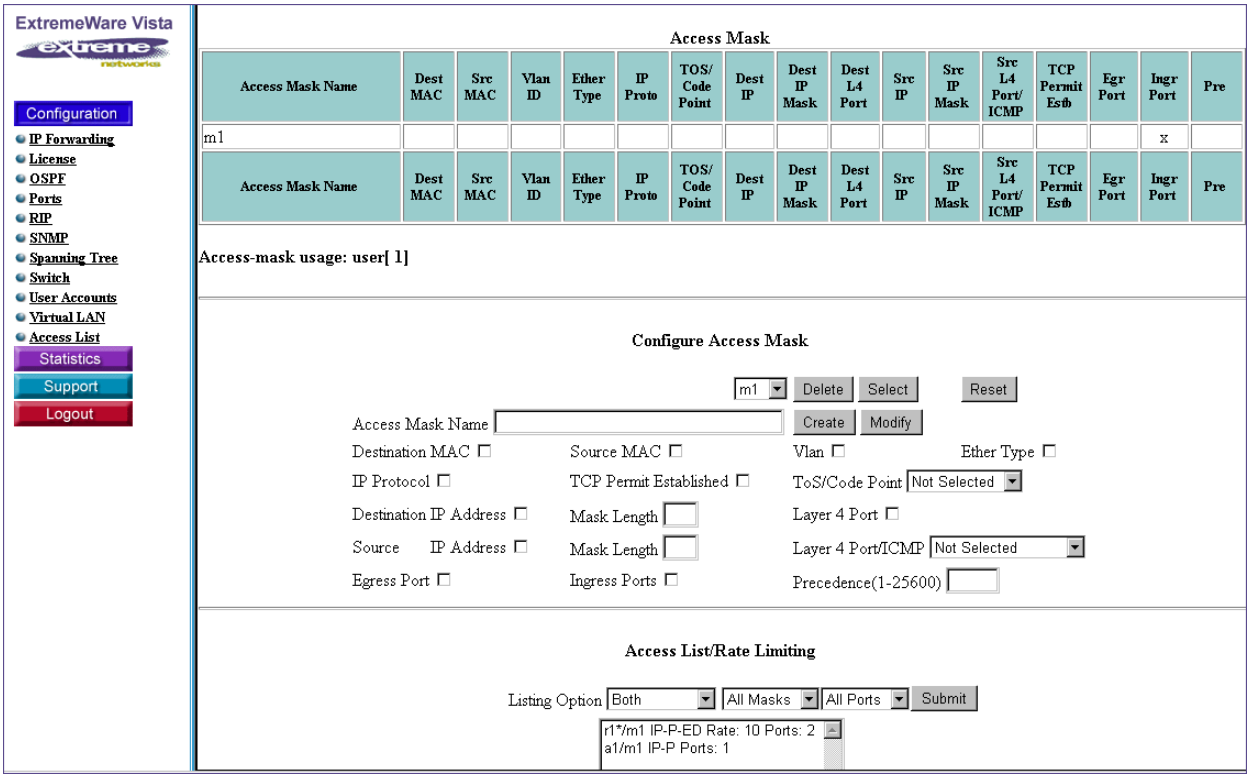

As Figure 75 shows, the next section of the window allows you to create, reset, modify or delete an access mask. Use the checkboxes to specify an option.

Rate Limiting

Like an access list, a rate limit includes a list of values to compare with the incoming packets and an action to take for packets that match. Additionally, a rate limit specifies an action to take when matching packets arrive at a rate above the limit you set. When you create a rate limit, you must specify a value for each of the fields that make up the access mask used by the list. Unlike an access list, a rate limit can only be applied to a single port. Each port has its own rate limit defined separately.

Each entry that makes up a rate limit contains a unique name and specifies a previously created access mask. As Figure 76 shows, the next section allows you configure an access list, a rate limit, or both. Use the pull-down menus to select the type of listing option, a mask, and the ports. Press the **Submit** button, when your configuration is complete.

Figure 76: Access List Configuration (2 of 2)

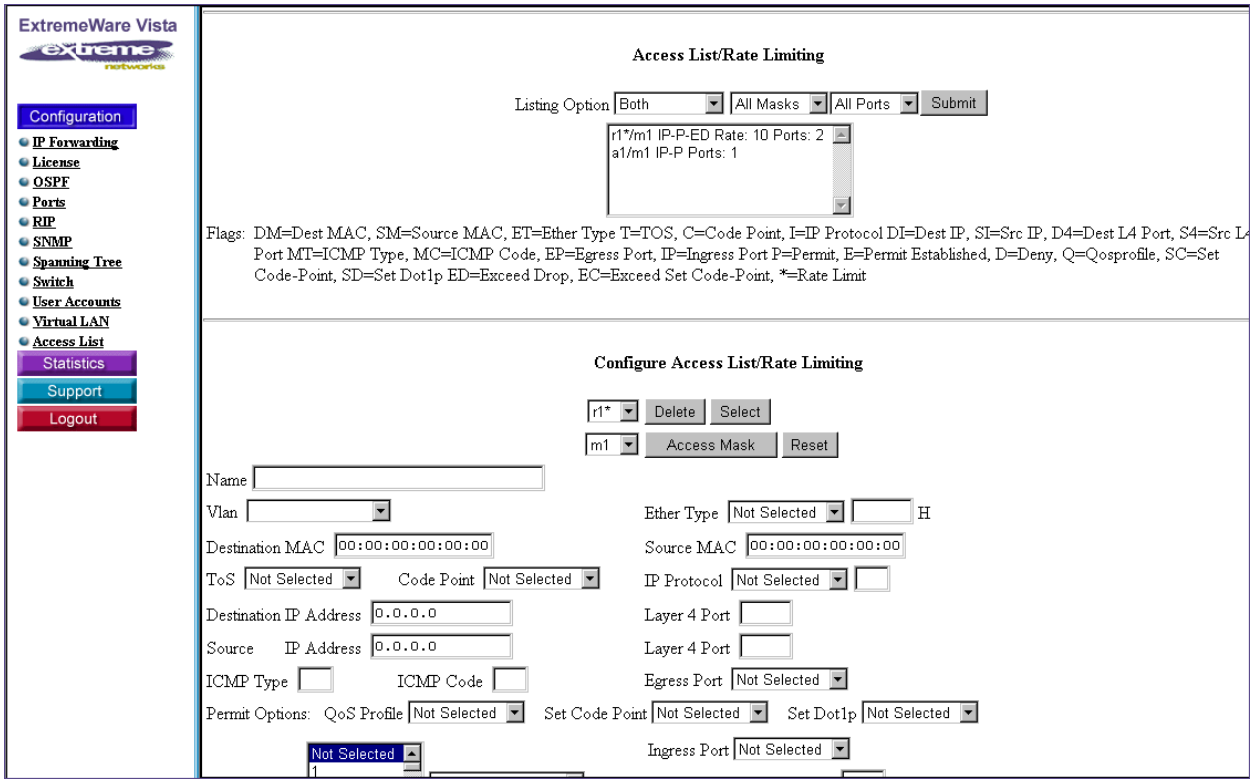

As shown in Figure 77, the final section of this window allows you to create, modify, or delete an access list. You can also create, modify or reset a rate limit. See the previous section for definitions of these fields.

Figure 77: Access List Configuration (4 of 4)

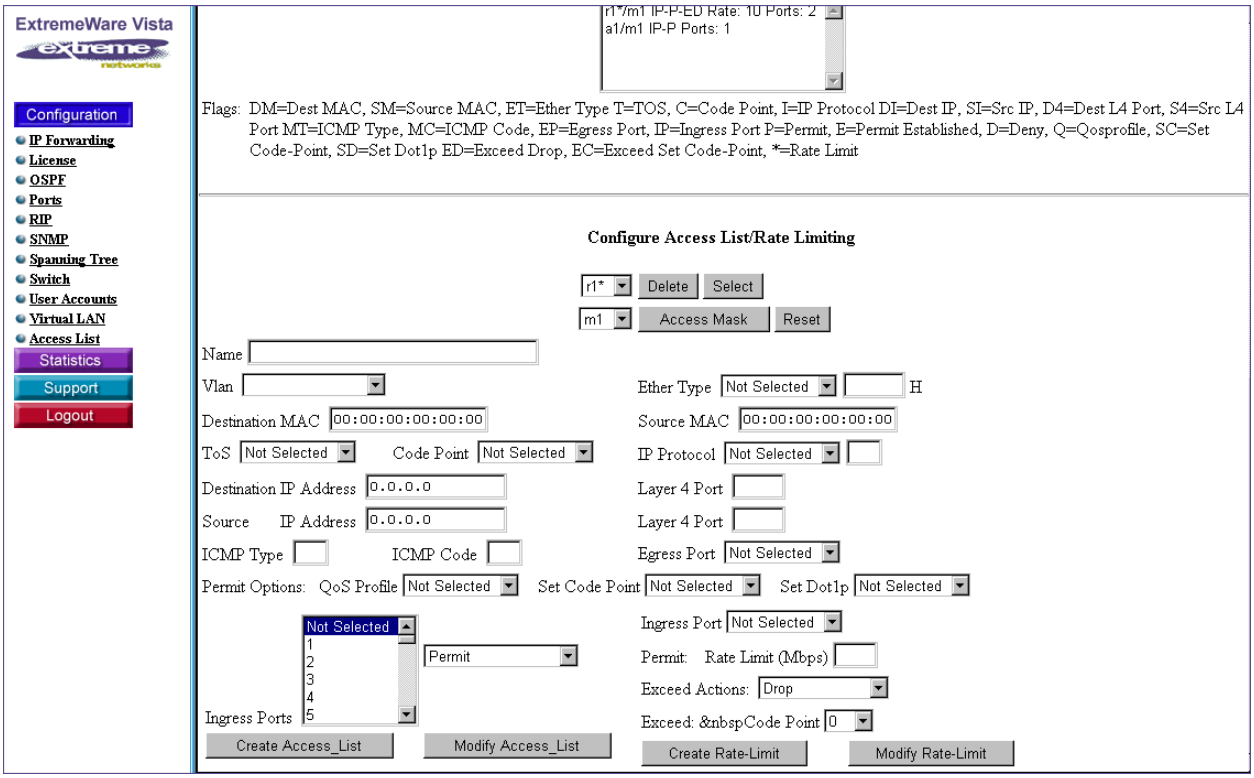

Reviewing ExtremeWare Vista Statistical Reports

ExtremeWare Vista offers a number of pre-formatted reports on the most frequently requested information. These statistical reports provide current information about the switch and its configuration.

To access the statistical reports, click **Statistics** in the task bar to reveal the submenu links. The following links appear in the submenu:

Event Log—Contains system event log entries

FDB—Contains Forwarding Database entries

IP ARP—Contains the entries in the IP Address Resolution Protocol (ARP) table

IP Configuration—Contains the global IP configuration statistics and router interface statistics

IP Route—Contains the IP Route table

IP Statistics—Contains global IP statistics

Ports—Contains the physical port statistics

Port Collisions—Contains Ethernet collision summary

Port Errors—Contains Ethernet port errors

Port Utilization—Contains link utilization information

RIP—Contains global RIP statistics and router interface statistics

Switch—Contains the hardware profile for the switch

Event Log

The System Even Log tracks all configuration and fault information pertaining to the device. Each entry in the log contains the following information:

- **Timestamp**—The timestamp records the month and day of the event, along with the time (hours, minutes, and seconds) in the form HH:MM:SS. If the event was caused by a user, the user name is also provided.
- **Fault level**—Describes the levels of importance that the system can assign to a fault. A fault level can either be classified as critical, warning, informational, or debug.

  By default, log entries that are assigned a critical or warning level remain in the log after a switch reboot. Issuing a clear log command does not remove these static entries.
- **Subsystem**—The subsystem refers to the specific functional area to which the error refers.

For additional information on system logging, see “Event Management System/Logging” on page 122.

Figure 78: Event Entries

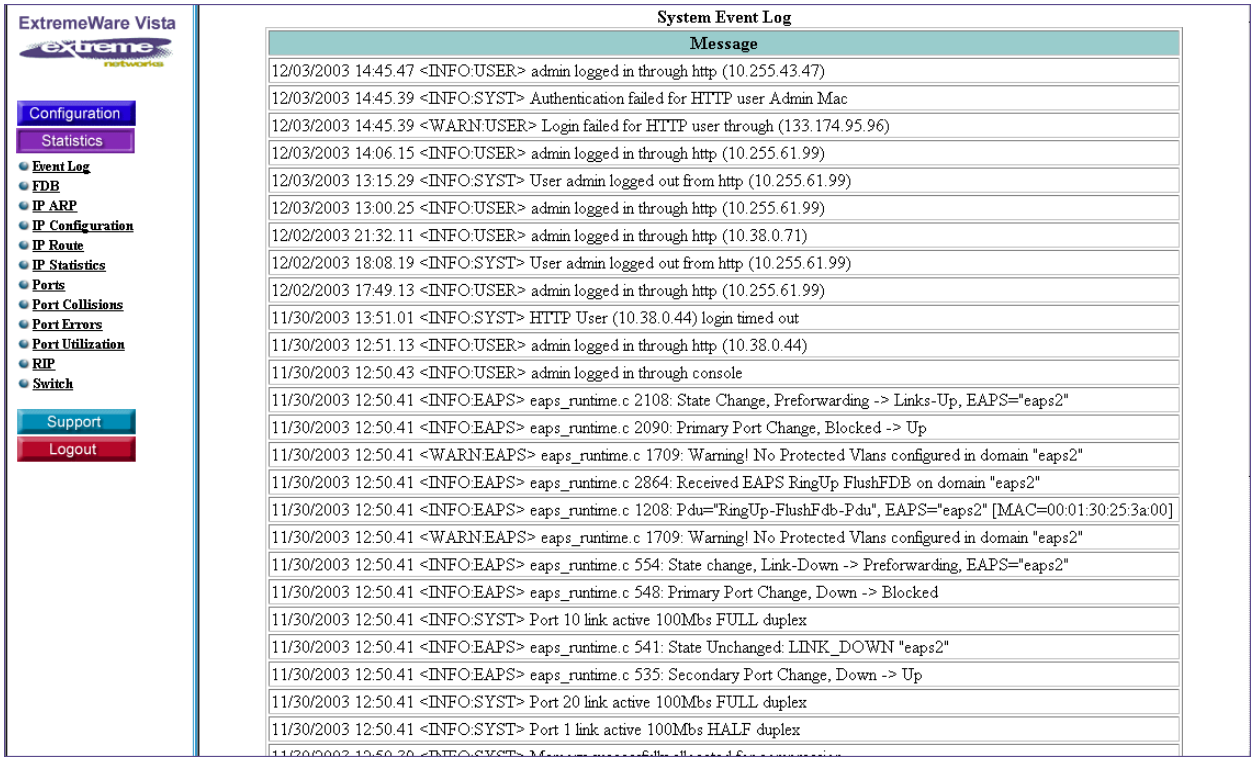

FDB

This window allows you to review the contents of the FDB table. It also gives summary information about the contents of the view and allows you tailor the view by various parameters.

The view of the FDB, as shown in Figure 79, consists of the following entries:

MAC Destination—MAC address of the device

VLAN—VLAN name and tag

Flags—Identifier for static (s) or dynamic (d)

Port List—The destination port or ports for the MAC address

Figure 79: FDB (1 of 2)

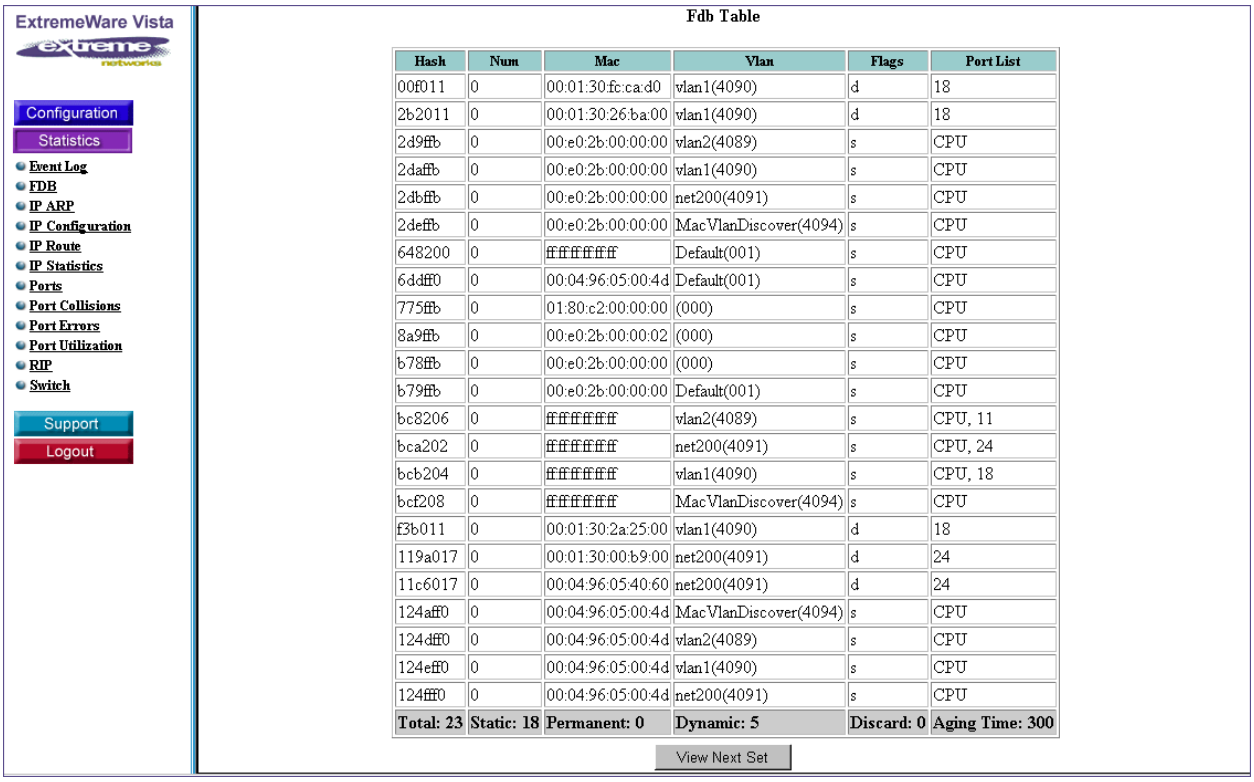

Hash	Num	Mac	Vlan	Flags	Port List
00f011	0	00:01:30:fc:ca:d0	vlan1(4090)	d	18
2b2011	0	00:01:30:26:ba:00	vlan1(4090)	d	18
2d9ffb	0	00:e0:2b:00:00:00	vlan2(4089)	s	CPU
2daffb	0	00:e0:2b:00:00:00	vlan1(4090)	s	CPU
2dbffb	0	00:e0:2b:00:00:00	net200(4091)	s	CPU
2deffb	0	00:e0:2b:00:00:00	MacVlanDiscover(4094)	s	CPU
648200	0	ffffffffff	Default(001)	s	CPU
6ddff0	0	00:04:96:05:00:4d	Default(001)	s	CPU
775ffb	0	01:80:c2:00:00:00	(000)	s	CPU
8a9ffb	0	00:e0:2b:00:00:02	(000)	s	CPU
b78ffb	0	00:e0:2b:00:00:00	(000)	s	CPU
b79ffb	0	00:e0:2b:00:00:00	Default(001)	s	CPU
bc8206	0	ffffffffff	vlan2(4089)	s	CPU, 11
bca202	0	ffffffffff	net200(4091)	s	CPU, 24
bcb204	0	ffffffffff	vlan1(4090)	s	CPU, 18
bcf208	0	ffffffffff	MacVlanDiscover(4094)	s	CPU
f3b011	0	00:01:30:2a:25:00	vlan1(4090)	d	18
119a017	0	00:01:30:00:b9:00	net200(4091)	d	24
11c6017	0	00:04:96:05:40:60	net200(4091)	d	24
124aff0	0	00:04:96:05:00:4d	MacVlanDiscover(4094)	s	CPU
124dff0	0	00:04:96:05:00:4d	vlan2(4089)	s	CPU
124eff0	0	00:04:96:05:00:4d	vlan1(4090)	s	CPU
124fff0	0	00:04:96:05:00:4d	net200(4091)	s	CPU
Total: 23	**Static: 18**	**Permanent: 0**	**Dynamic: 5**	**Discard: 0**	**Aging Time: 300**

Summary information is located at the bottom of the view. The summary information contains the:

Total—Total number of entries in this database view

Static—Number of static entries in this view

Permanent—Number of permanent entries in this view

Dynamic—Number of dynamic entries in this view

Discarded—Number of entries discarded

Aging Time—The current time setting for removing entries from the FDB

The View Options allow you to filter and restrict the amount of information presented in the FDB view.

Figure 80: FDB (2 of 2)

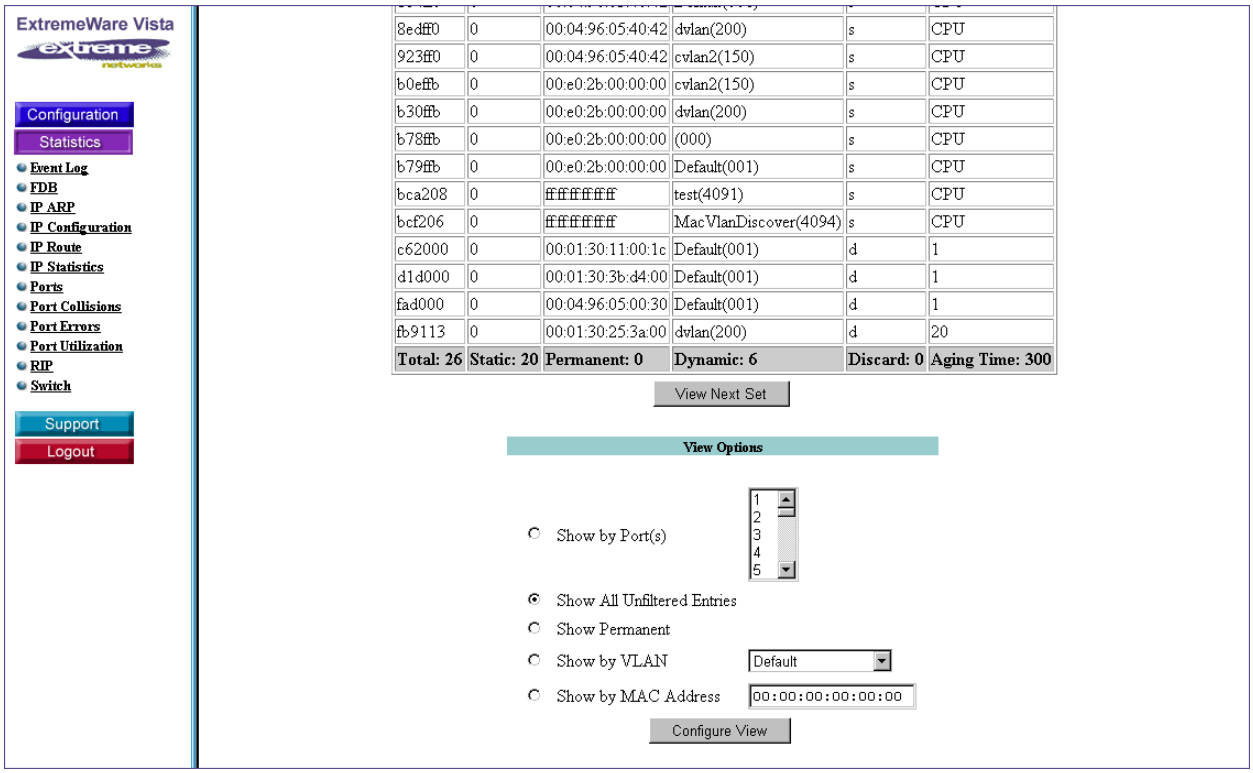

For further information about the FDB, see “Overview of the FDB” on page 99.

IP ARP

Use the IP ARP to find the MAC address associated with an IP address.

The IP ARP table contains the following fields:

Destination—The destination IP address

MAC Address—The MAC address associated with the IP address

Age—The age of the entry

Flags—Identifier for static entry (m), proxy ARP (p), and trailers requested (t)

Static—Either yes for a static entry or no for dynamic

VLAN—VLAN name

VLAN ID

Figure 81: IP ARP Table

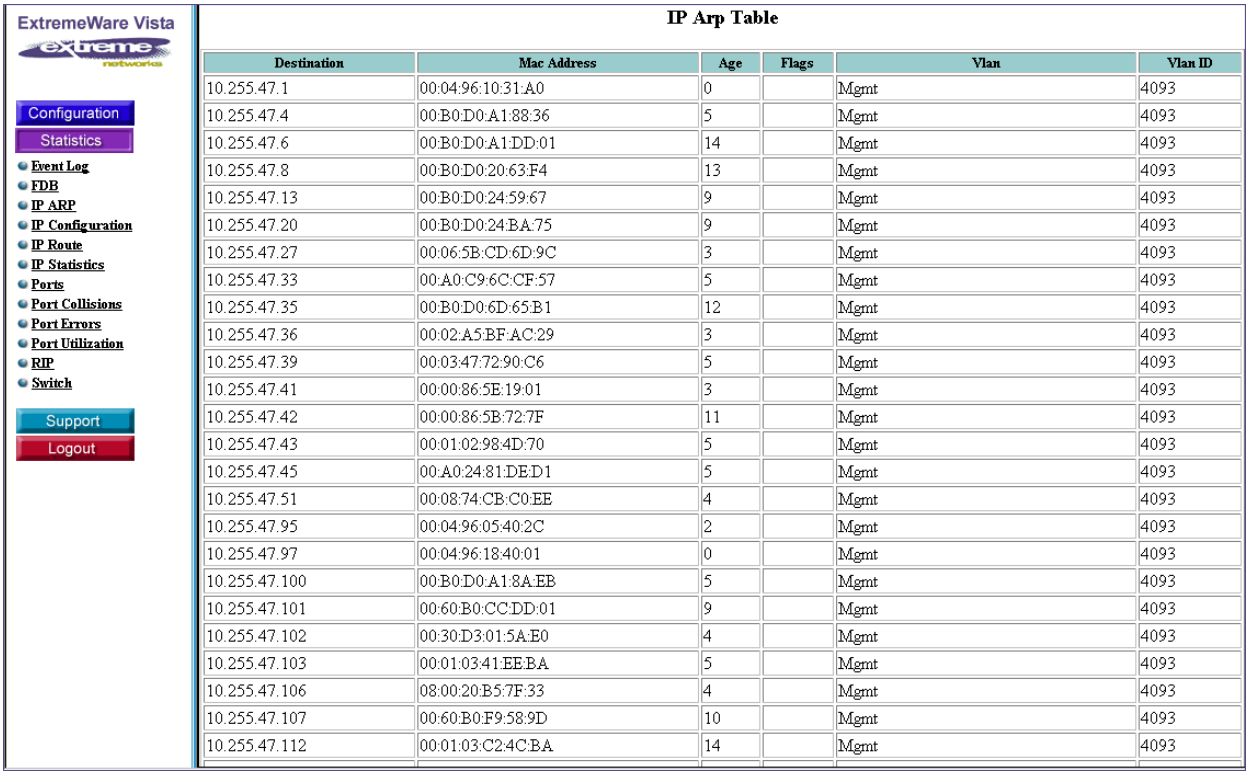

IP Arp Table

Destination	Mac Address	Age	Flags	Vlan	Vlan ID
10.255.47.1	00:04:96:10:31:A0	0		Mgmt	4093
10.255.47.4	00:B0:D0:A1:88:36	5		Mgmt	4093
10.255.47.6	00:B0:D0:A1:DD:01	14		Mgmt	4093
10.255.47.8	00:B0:D0:20:63:F4	13		Mgmt	4093
10.255.47.13	00:B0:D0:24:59:67	9		Mgmt	4093
10.255.47.20	00:B0:D0:24:BA:75	9		Mgmt	4093
10.255.47.27	00:06:5B:CD:6D:9C	3		Mgmt	4093
10.255.47.33	00:A0:C9:6C:CF:57	5		Mgmt	4093
10.255.47.35	00:B0:D0:6D:65:B1	12		Mgmt	4093
10.255.47.36	00:02:A5:BF:AC:29	3		Mgmt	4093
10.255.47.39	00:03:47:72:90:C6	5		Mgmt	4093
10.255.47.41	00:00:86:5E:19:01	3		Mgmt	4093
10.255.47.42	00:00:86:5B:72:7F	11		Mgmt	4093
10.255.47.43	00:01:02:98:4D:70	5		Mgmt	4093
10.255.47.45	00:A0:24:81:DE:D1	5		Mgmt	4093
10.255.47.51	00:08:74:CB:C0:EE	4		Mgmt	4093
10.255.47.95	00:04:96:05:40:2C	2		Mgmt	4093
10.255.47.97	00:04:96:18:40:01	0		Mgmt	4093
10.255.47.100	00:B0:D0:A1:8A:EB	5		Mgmt	4093
10.255.47.101	00:60:B0:CC:DD:01	9		Mgmt	4093
10.255.47.102	00:30:D3:01:5A:E0	4		Mgmt	4093
10.255.47.103	00:01:03:41:EE:BA	5		Mgmt	4093
10.255.47.106	08:00:20:B5:7F:33	4		Mgmt	4093
10.255.47.107	00:60:B0:F9:58:9D	10		Mgmt	4093
10.255.47.112	00:01:03:C2:4C:BA	14		Mgmt	4093

IP Configuration

In this window you can review two different tables containing IP configuration information. The Global IP Configuration Statistics table provides IP settings and summary statistics for the entire switch. The Router Interface table provides details on each VLAN. Both tables are shown in Figure 82.

Global IP Configuration Statistics

This table contains the following fields:

IP Routing—Indicates whether IP forwarding is either enabled or disabled on the switch. The default setting for IP forwarding is disabled.

Ipmc Routing— Indicates whether IP multicast forwarding is enabled or disabled on the switch. This setting is either enabled or disabled.

Use Redirects—Indicates whether the switch can modify the route table information when an ICMP redirect message is received. This option applies to the switch when it is not configured for routing. This setting is either enabled or disabled; the default setting is disabled.

IGMP—Internet Group Management Protocol (IGMP) allows network hosts to report the multicast group membership to the switch. This setting is either enabled or disabled.

RIP—Routing Information Protocol (RIP) is either enabled or disabled.

IRDP—ICMP Router Discovery Protocol (IRDP) shows the generation of ICMP router advertisement messages on one or all VLANs. The setting is either enabled or disabled; the default setting is enabled.

OSPF—The OSPF routing protocol for the switch. The setting is either enabled or disabled.

Advertisement Address—The destination address of the router advertisement messages.

Maximum Interval—The maximum time between router advertisements. The default setting is 600 seconds.

Minimum Interval—The minimum amount of time between router advertisements. The default setting is 450 seconds.

Lifetime—The client aging timer setting, the default is 1,800 seconds.

Preference—The preference level of the router. An IRDP client always uses the router with the highest preference level. The default setting is 0.

Bootp Relay—The BOOTP relay service on the switch. The setting is either enabled or disabled; the default is disabled.

Figure 82: IP Configuration Statistics

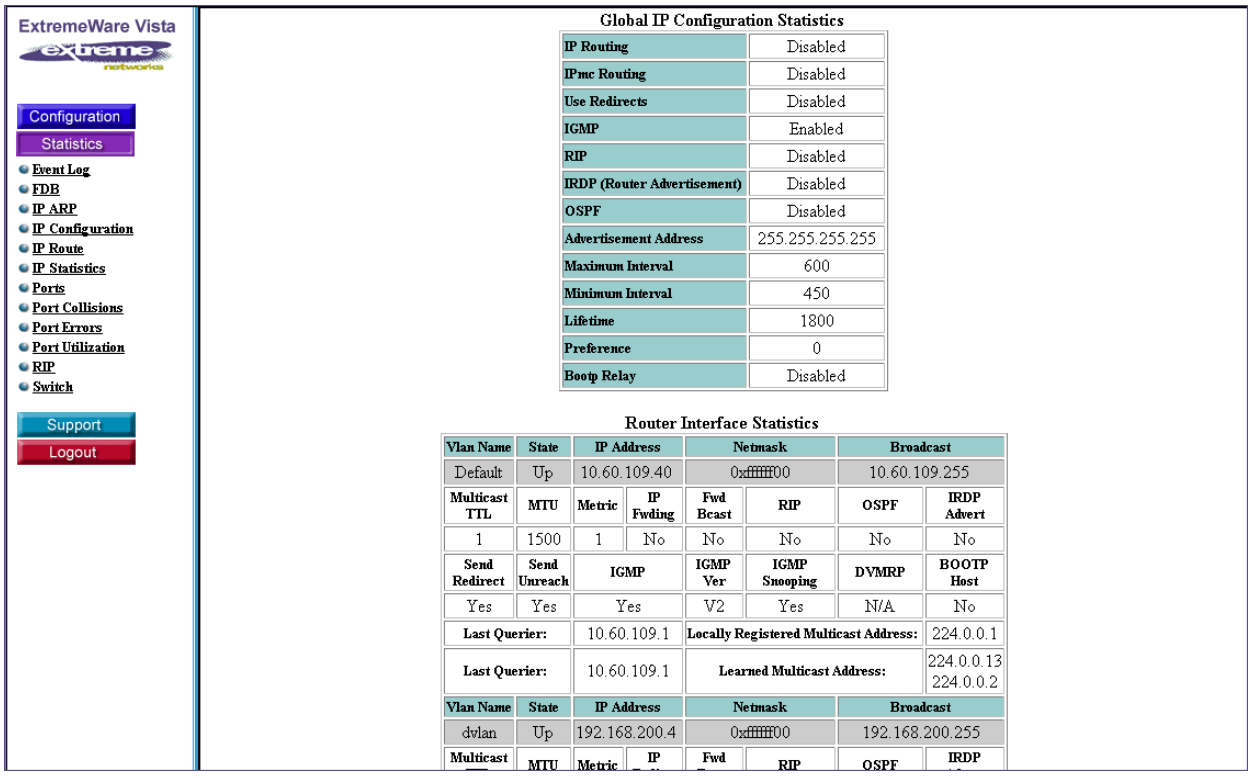

Router Interface Statistics

The Router Interface Statistics table gives the details of individual VLANs. It contains the following fields:

VLAN name

State—up or down

IP Address—in dotted-quad notation

Netmask

Broadcast—The broadcast address in dotted-quad notation

Multicast TTL—The multicast time-to-live

MTU—Maximum Transmission Unit (MTU) size

Metric—The hop count to the destination address

IP Forwarding—IP forwarding on this interface is enabled or disabled

Fwd Bcast—The hardware forwarding of subnet-directed broadcast IP packets is enabled or disabled

RIP—RIP is enabled or disabled on this interface

OSPF—OSPF is enabled or disabled on this interface

IDRP—IDRP is enabled or disabled on this interface

Send Redirect—Allows or disallows the interface to modify the route table information when an ICMP redirect message is received

Send Unreach—Allows or disallows the interface to generate an ICMP port unreachable messages (type 3, code 3) when a TPC or UDP request is made to the switch, and no application is waiting for the request, or access policy denies the request.

IGMP—IGMP is enabled or disabled on this interface

IGMP Ver—The version of IGMP running on the interface

IGMP Snooping—Enable or disable of IGMP Snooping

BOOTP Host—Indicates whether BOOTP is enabled on this VLAN or not

Last Querier—The address of the querier

Locally Registered Multicast Address

Learned Multicast Address

IP Route

This window contains the statistics for the IP routing table. The Summit 400 exchanges routing information with other routers and switches on the network using either the RIP or the OSPF protocol. The Summit 400 dynamically builds and maintains the routing table, and determines the best path for each of its routes.

The IP route table contains the following fields:

Destination—The destination address

Gateway—The gateway address

Mtr—The cost metric

Flags—For example, U for ub; G for gateway; and U for unicast

Use—The number of times the entry is used

VLAN—VLAN name

Origin—Route origin. One of the following:

- direct
- blackhole
- static
- ICMP
- OSPFIntra
- OSPFInter
- RIP
- OSPFExtern1
- OSPFExtern2
- BOOTP

As shown in Figure 83, you can also use the View Options to restrict different aspects of the view. For more information on IP routing, see “Populating the Routing Table” on page 217.

Figure 83: IP Route Table

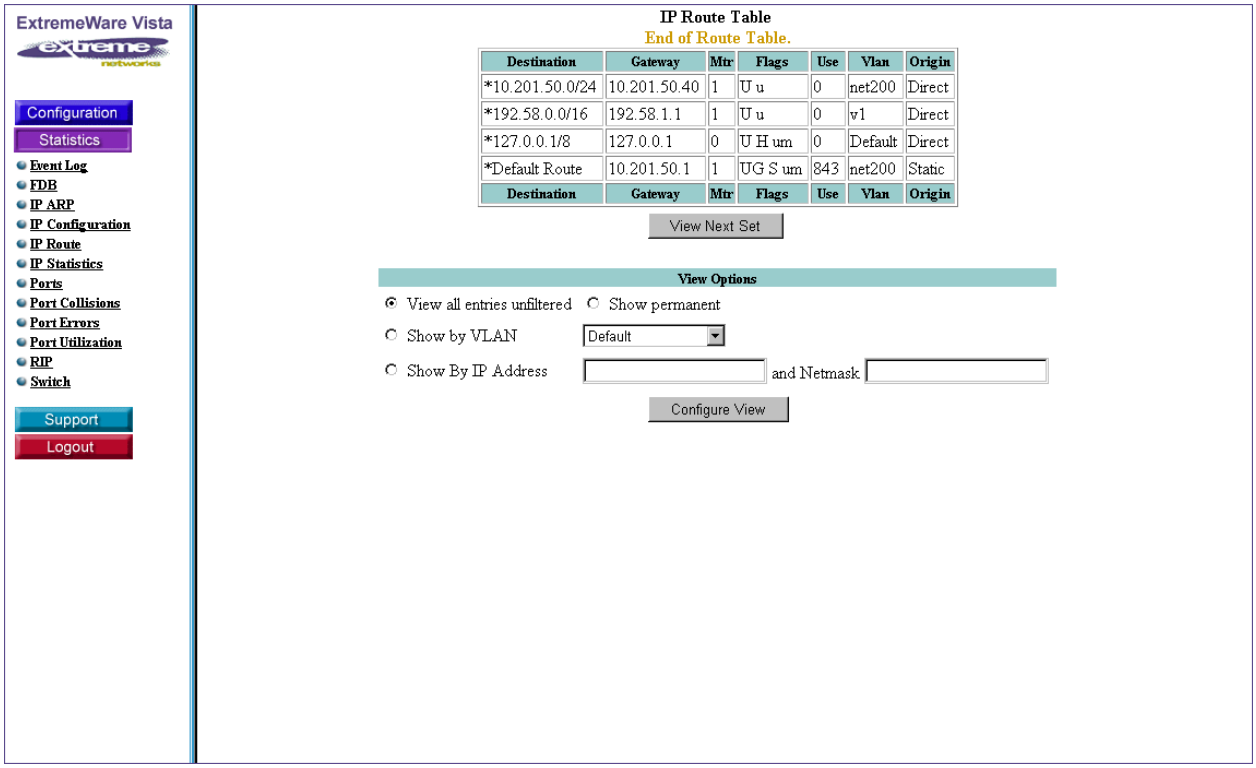

IP Statistics

This window provides ICMP error reporting statistics and error counts from the switch as a whole, and also on individual interfaces. For information about error counts:

- Across the whole switch, see “Global IP Statistics”
- On an interface, see “Global ICMP Statistics” on page 291
- Across VLANs, see “Global ICMP Statistics” on page 292

Global IP Statistics

The Global IP Statistics report IP traffic flow through the switch. As shown at the top of Figure 84, these statistics are grouped into four logical groups:

- Inbound traffic
- Outbound traffic
- Bad packets received
- Other types of errors

Figure 84: Global IP Statistics

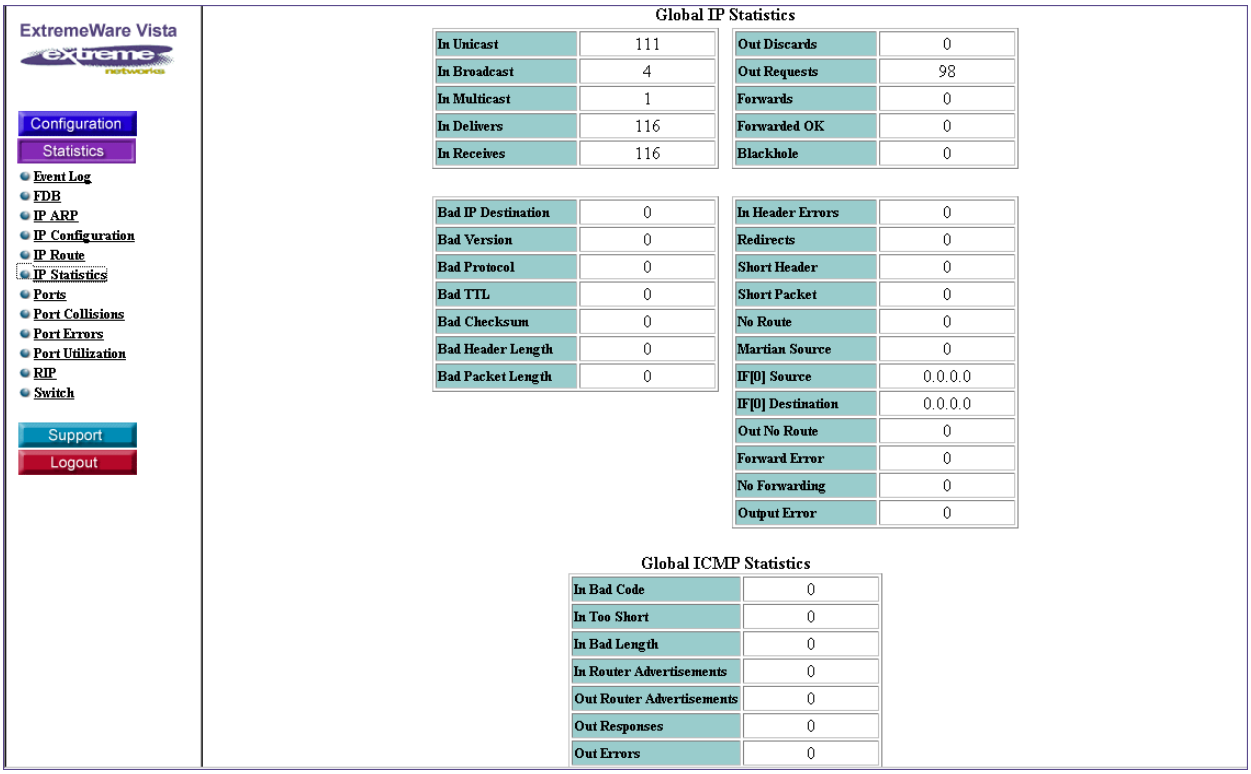

Global ICMP Statistics

ICMP provides error reporting, flow control and first-hop gateway redirection. As shown in Figure 85, the Global ICMP Statistics table provides information about error counts found in the following areas:

- In Bad Code
- In Too Short
- In Bad Length
- In Router Advertisements
- Out Router Advertisements

- Out Responses
- Out Errors
- Bad Checksums

Figure 85: Global ICMP Statistics

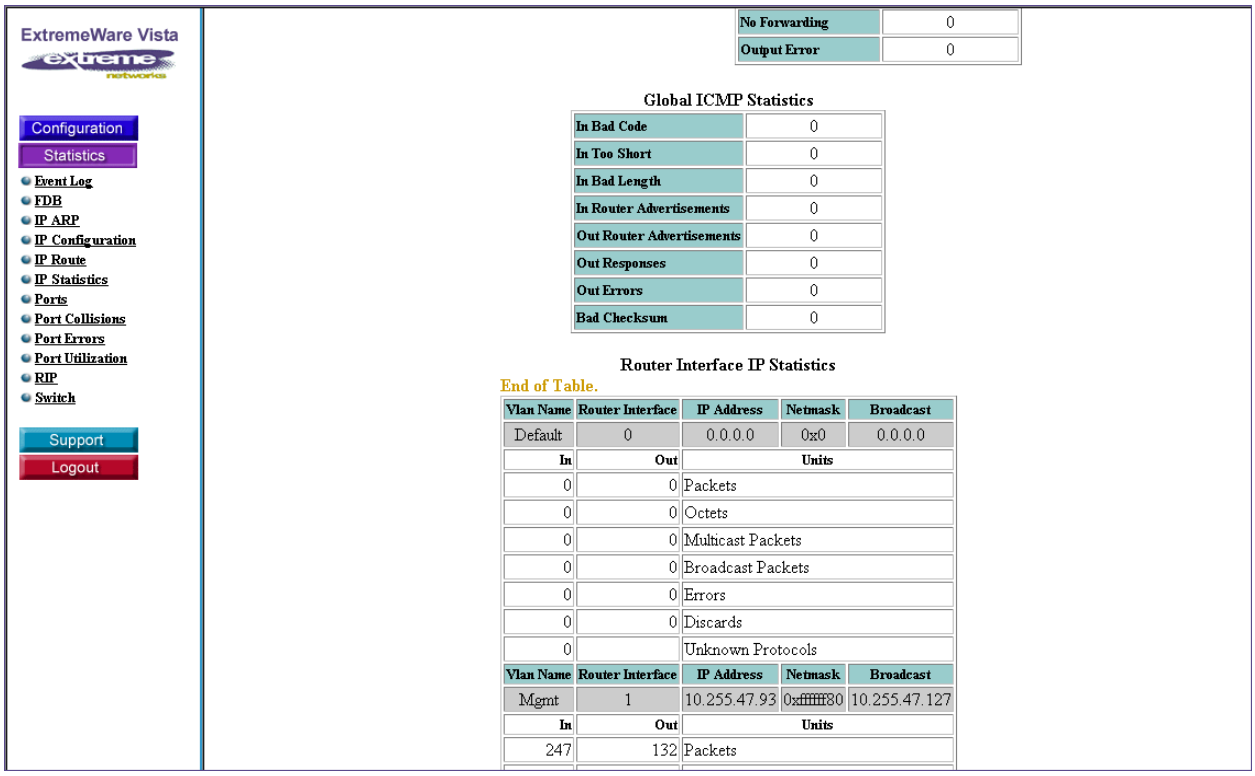

Router Interface IP Statistics

The Router Interface IP Statistics give detailed traffic details at the VLAN level, as shown in Figure 86. For each interface the table provides:

- VLAN name
- Interface ID
- IP Address
- Netmask
- Broadcast Address
- Amount in and out of the switch for the following units: packets, octets, multicast packets, broadcast packets, errors, discards, and unknown protocols

Figure 86: Router Interface IP Statistics

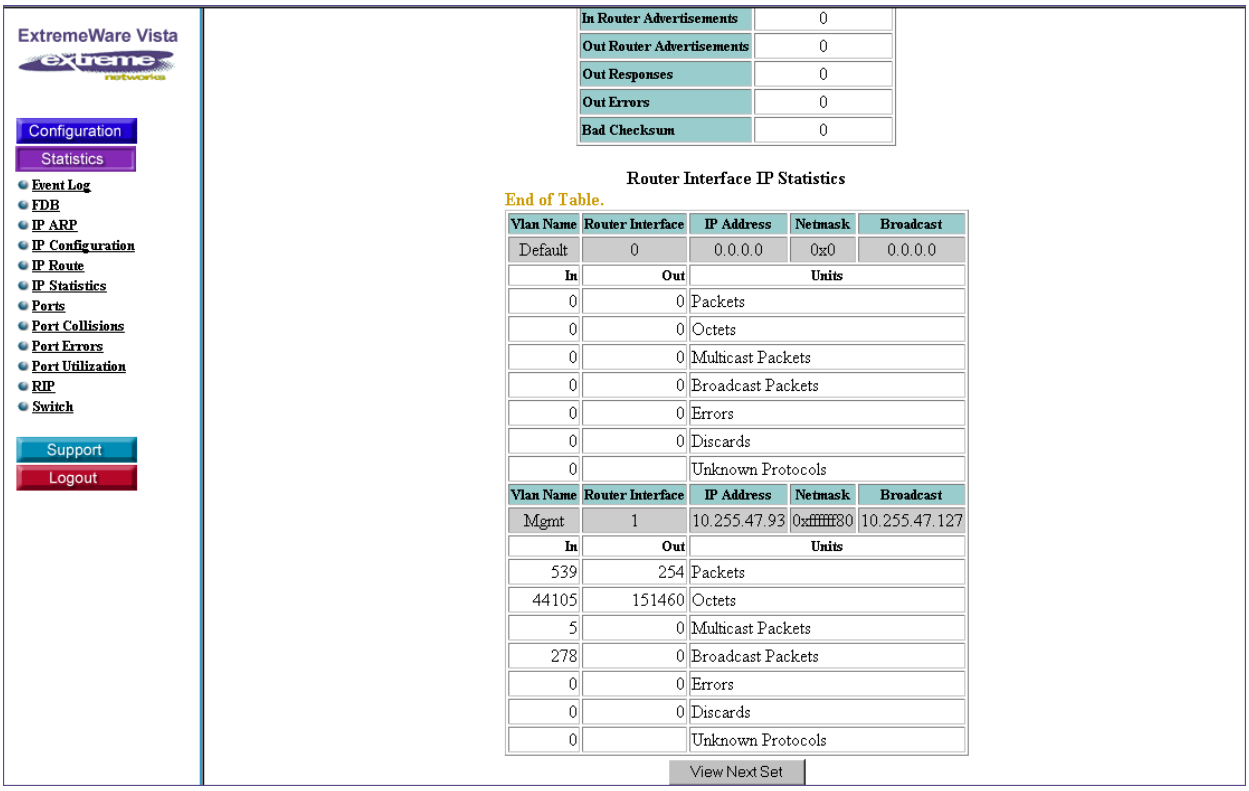

Ports

This window provides information about active ports as reported by the Summit 400 hardware. As shown in Figure 87, the report consists of the following fields:

Port Number

Port Speed

Link State

Received Packet Count

Transmitted Packet Count

Received Byte Count

Transmitted Byte Count

Collisions

Figure 87: Physical Port Statistics

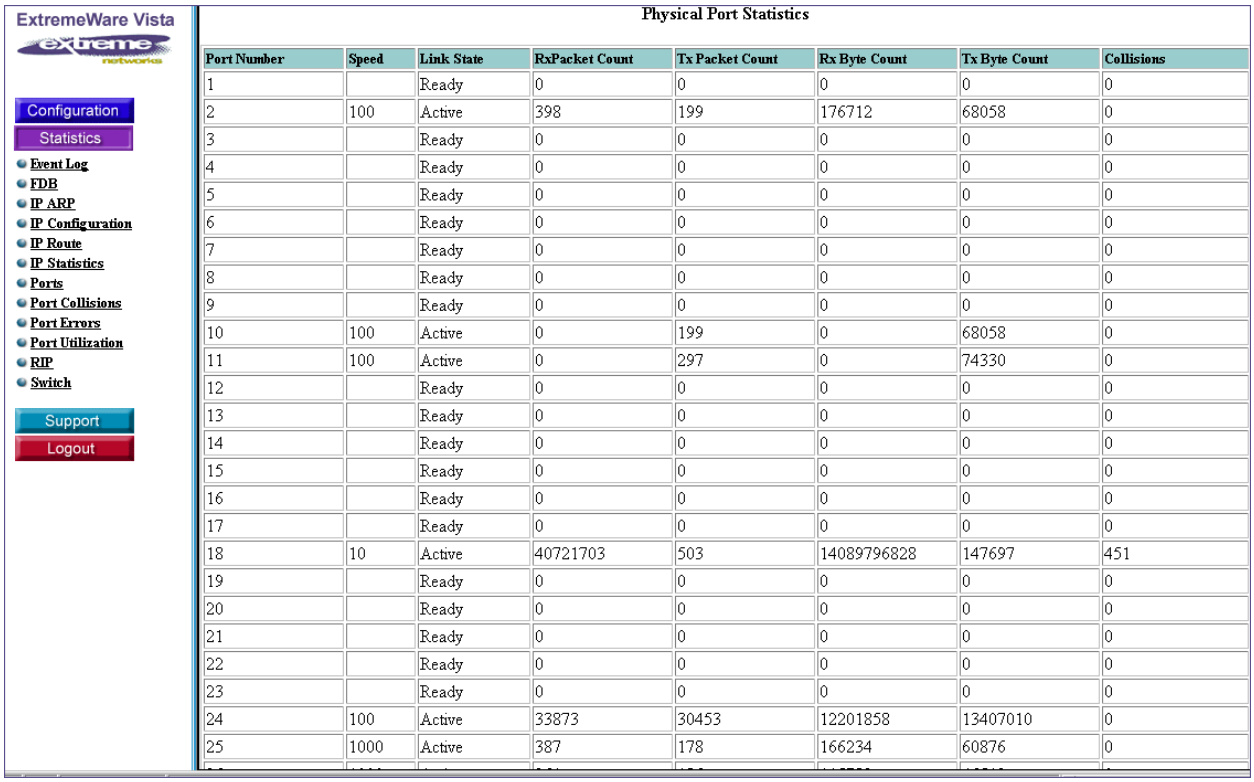

Physical Port Statistics

Port Number	Speed	Link State	RxPacket Count	Tx Packet Count	Rx Byte Count	Tx Byte Count	Collisions
1		Ready	0	0	0	0	0
2	100	Active	398	199	176712	68058	0
3		Ready	0	0	0	0	0
4		Ready	0	0	0	0	0
5		Ready	0	0	0	0	0
6		Ready	0	0	0	0	0
7		Ready	0	0	0	0	0
8		Ready	0	0	0	0	0
9		Ready	0	0	0	0	0
10	100	Active	0	199	0	68058	0
11	100	Active	0	297	0	74330	0
12		Ready	0	0	0	0	0
13		Ready	0	0	0	0	0
14		Ready	0	0	0	0	0
15		Ready	0	0	0	0	0
16		Ready	0	0	0	0	0
17		Ready	0	0	0	0	0
18	10	Active	40721703	503	14089796828	147697	451
19		Ready	0	0	0	0	0
20		Ready	0	0	0	0	0
21		Ready	0	0	0	0	0
22		Ready	0	0	0	0	0
23		Ready	0	0	0	0	0
24	100	Active	33873	30453	12201858	13407010	0
25	1000	Active	387	178	166234	60876	0

Port Collisions

This window provides information about Ethernet collisions that occur when the port is operating in half-duplex mode. An example of this window is shown in Figure 88.

Figure 88: Port Collisions

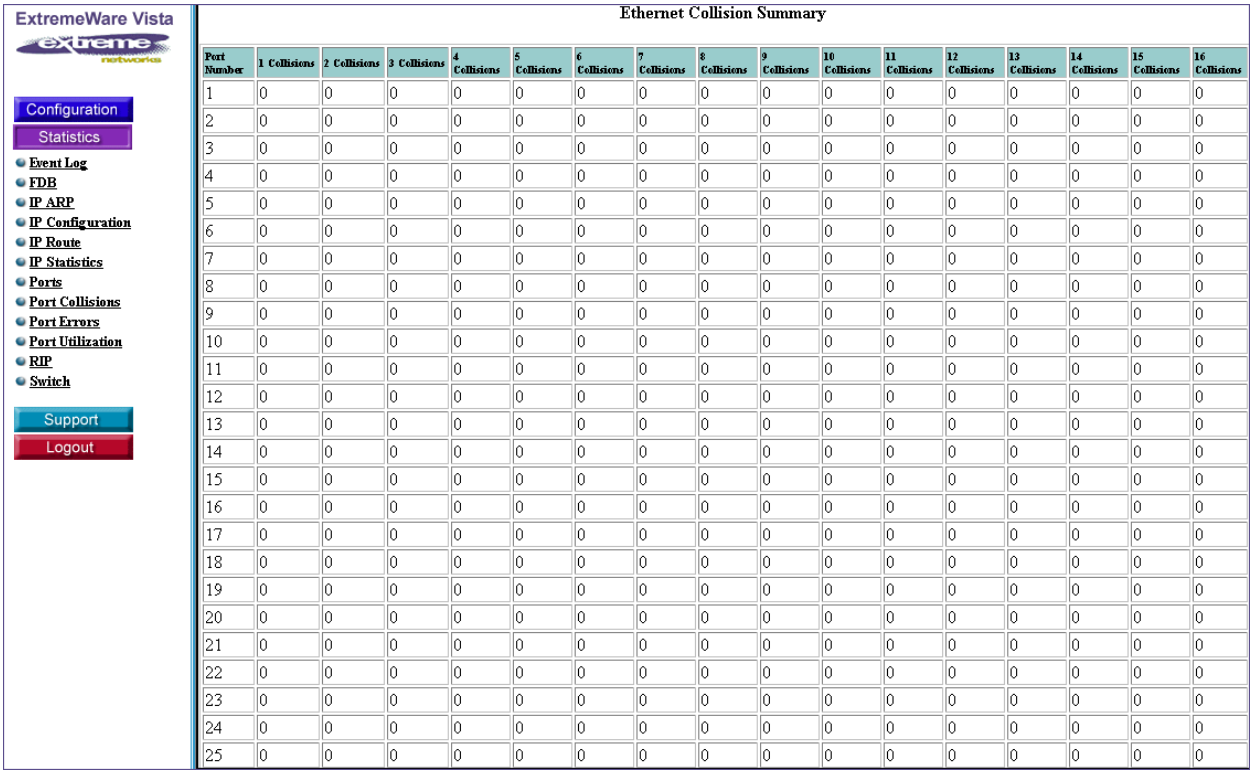

Ethernet Collision Summary

Port Number	1 Collisions	2 Collisions	3 Collisions	4 Collisions	5 Collisions	6 Collisions	7 Collisions	8 Collisions	9 Collisions	10 Collisions	11 Collisions	12 Collisions	13 Collisions	14 Collisions	15 Collisions	16 Collisions
1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
6	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
7	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
8	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
9	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
10	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
11	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
12	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
13	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
14	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
15	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
16	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
17	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
18	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
19	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
20	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
21	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
22	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
23	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
24	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
25	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Port Errors

In this window, you can review Ethernet link errors. As shown in Figure 89, the table reflects the following information for each active port:

- Link State
- Rx Lost
- Rx Bad Cyclic Redundancy Check (CRC)
- Rx Undersize
- Rx Oversize
- Rx Fragments
- Rx Jabber
- Rx Alignment
- Tx Errored
- Tx Deferred
- Tx Late Collisions

Figure 89: Ethernet Port Errors

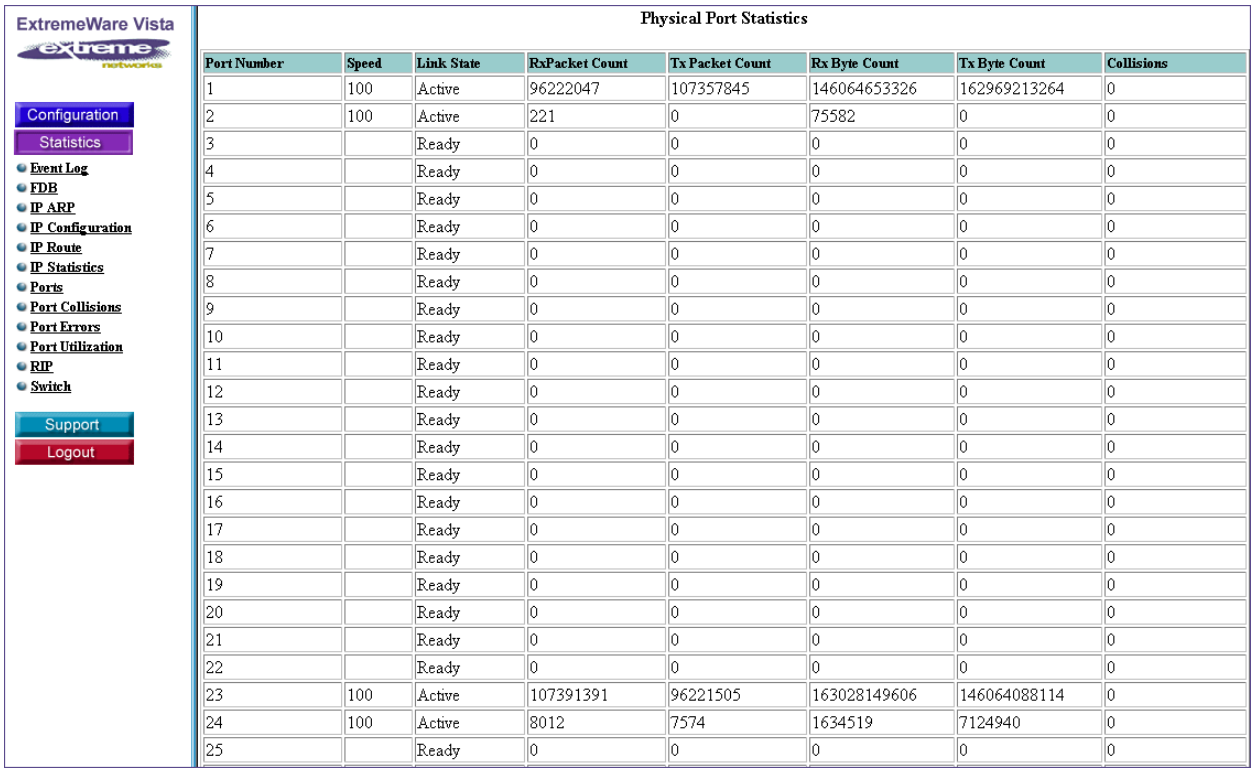

Physical Port Statistics

Port Number	Speed	Link State	RxPacket Count	Tx Packet Count	Rx Byte Count	Tx Byte Count	Collisions
1	100	Active	96222047	107357845	146064653326	162969213264	0
2	100	Active	221	0	75582	0	0
3		Ready	0	0	0	0	0
4		Ready	0	0	0	0	0
5		Ready	0	0	0	0	0
6		Ready	0	0	0	0	0
7		Ready	0	0	0	0	0
8		Ready	0	0	0	0	0
9		Ready	0	0	0	0	0
10		Ready	0	0	0	0	0
11		Ready	0	0	0	0	0
12		Ready	0	0	0	0	0
13		Ready	0	0	0	0	0
14		Ready	0	0	0	0	0
15		Ready	0	0	0	0	0
16		Ready	0	0	0	0	0
17		Ready	0	0	0	0	0
18		Ready	0	0	0	0	0
19		Ready	0	0	0	0	0
20		Ready	0	0	0	0	0
21		Ready	0	0	0	0	0
22		Ready	0	0	0	0	0
23	100	Active	107391391	96221505	163028149606	146064088114	0
24	100	Active	8012	7574	1634519	7124940	0
25		Ready	0	0	0	0	0

Port Utilization

This window shows port utilization. As shown in Figure 90, the report fields are as follows:

Port Number

Speed—Configured port speed, either 10, 100, 1000, or auto

Link Status—Either active (A) or ready (R)

Rx Pkt/Sec—Received packets rate

Peak Rx Pkt/Sec—Peak received packet rate

Tx Pkt/Sec—Transmission packet rate

Peak Tx Pkt/Sec—Peak packet rate transmitted

Rx Byte/Sec—Received byte rate

Peak Rx Byte/Sec—Peak received bytes rate

Tx Byte/Sec—Transmission byte rate

Peak Tx Byte/Sec—Peak transmission byte rate

Bandwidth—Bandwidth utilization

Peak Bandwidth—Peak bandwidth utilization

Figure 90: Utilization Averages

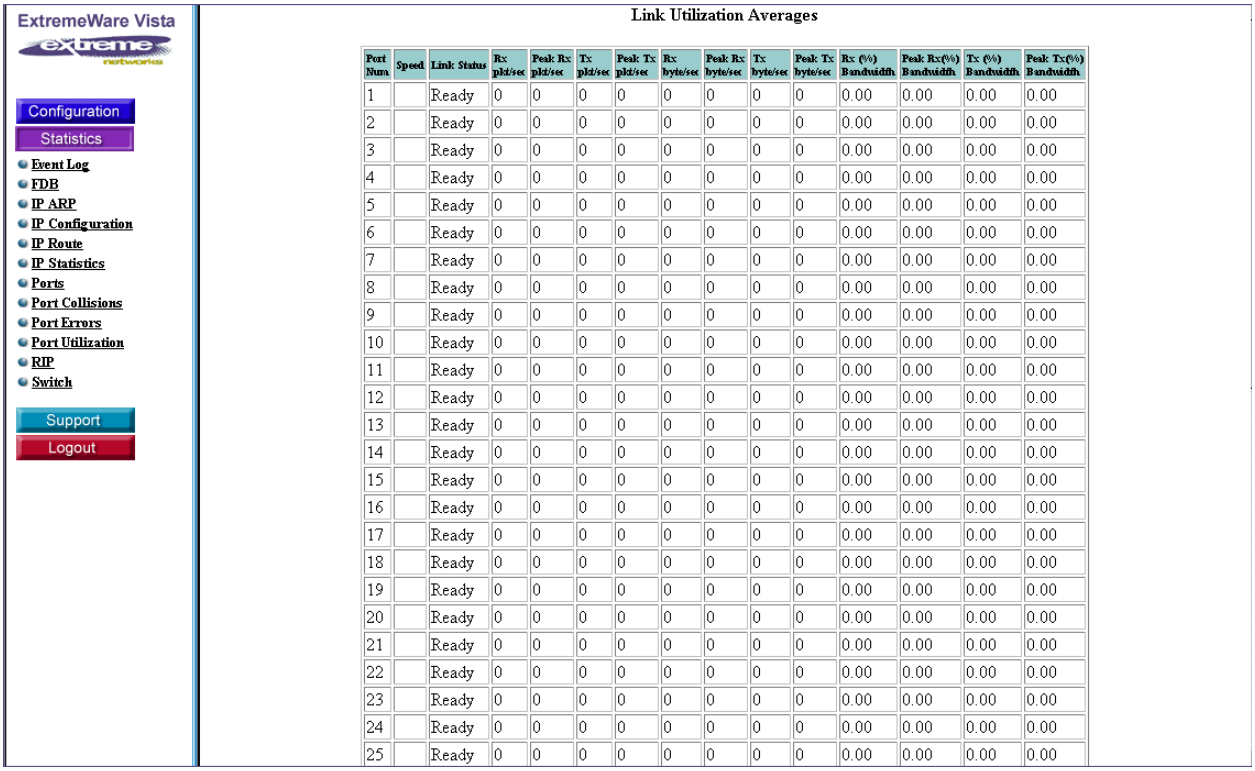

Port Num	Speed	Link Status	Rx pkt/sec	Peak Rx pkt/sec	Tx pkt/sec	Peak Tx pkt/sec	Rx byte/sec	Peak Rx byte/sec	Tx byte/sec	Peak Tx byte/sec	Rx (%) Bandwidth	Peak Rx(%) Bandwidth	Tx (%) Bandwidth	Peak Tx(%) Bandwidth
1		Ready	0	0	0	0	0	0	0	0	0.00	0.00	0.00	0.00
2		Ready	0	0	0	0	0	0	0	0	0.00	0.00	0.00	0.00
3		Ready	0	0	0	0	0	0	0	0	0.00	0.00	0.00	0.00
4		Ready	0	0	0	0	0	0	0	0	0.00	0.00	0.00	0.00
5		Ready	0	0	0	0	0	0	0	0	0.00	0.00	0.00	0.00
6		Ready	0	0	0	0	0	0	0	0	0.00	0.00	0.00	0.00
7		Ready	0	0	0	0	0	0	0	0	0.00	0.00	0.00	0.00
8		Ready	0	0	0	0	0	0	0	0	0.00	0.00	0.00	0.00
9		Ready	0	0	0	0	0	0	0	0	0.00	0.00	0.00	0.00
10		Ready	0	0	0	0	0	0	0	0	0.00	0.00	0.00	0.00
11		Ready	0	0	0	0	0	0	0	0	0.00	0.00	0.00	0.00
12		Ready	0	0	0	0	0	0	0	0	0.00	0.00	0.00	0.00
13		Ready	0	0	0	0	0	0	0	0	0.00	0.00	0.00	0.00
14		Ready	0	0	0	0	0	0	0	0	0.00	0.00	0.00	0.00
15		Ready	0	0	0	0	0	0	0	0	0.00	0.00	0.00	0.00
16		Ready	0	0	0	0	0	0	0	0	0.00	0.00	0.00	0.00
17		Ready	0	0	0	0	0	0	0	0	0.00	0.00	0.00	0.00
18		Ready	0	0	0	0	0	0	0	0	0.00	0.00	0.00	0.00
19		Ready	0	0	0	0	0	0	0	0	0.00	0.00	0.00	0.00
20		Ready	0	0	0	0	0	0	0	0	0.00	0.00	0.00	0.00
21		Ready	0	0	0	0	0	0	0	0	0.00	0.00	0.00	0.00
22		Ready	0	0	0	0	0	0	0	0	0.00	0.00	0.00	0.00
23		Ready	0	0	0	0	0	0	0	0	0.00	0.00	0.00	0.00
24		Ready	0	0	0	0	0	0	0	0	0.00	0.00	0.00	0.00
25		Ready	0	0	0	0	0	0	0	0	0.00	0.00	0.00	0.00

RIP

This window provides statistics about the Routing Information Protocol (RIP) both at the global (switch level) and at the interface level. At the switch level, the Global Routing Information Protocol Statistics table shows the number of route changes and the number of queries. As shown in Figure 91, at the interface level, the Router Interface Statistics table shows the following fields:

VLAN Name

Authentication—Yes for enabled, no for disabled on the interface

Rcvd Pkts—Received RIP packets

Sent Pkts—Sent RIP packets

Rcvd Bad Pkts—Received bad RIP packets

Rcvd Bad Routes—Received bad routes

Sent Trig Updts—Sent triggered updates

Peer

Age (sec)—Age in seconds

Version—RIP version

Bad Pkts—Bad Packets

Bad Routes

Figure 91: RIP Statistics

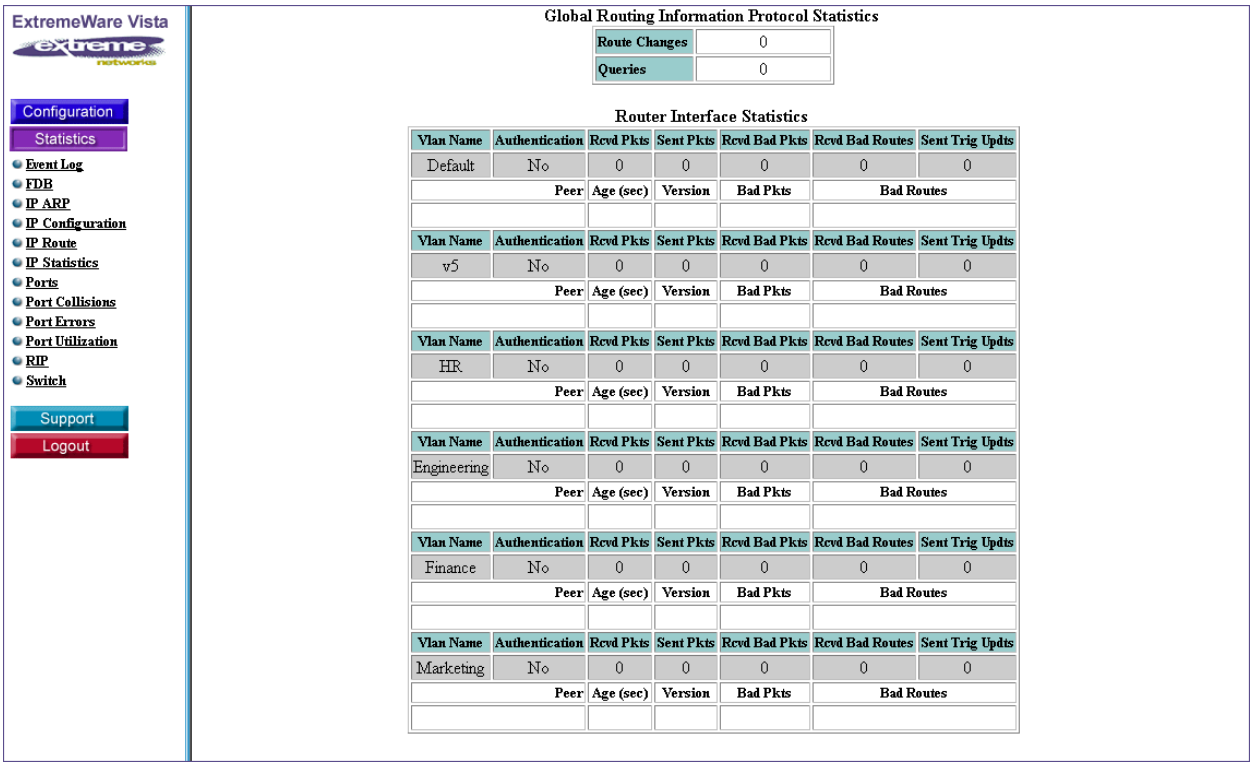

Switch

Use this window to locate hardware status information. As shown in Figure 92, the Hardware Status table provides data about the following areas:

System Name—Summit 400-48t

MAC Address—MAC address of the device

Software Image Selected—Primary or secondary image and version number of the image

Software Image Booted—Actual image running

Configuration Selected—Either primary or secondary

Configuration Booted—Either primary or secondary

Primary Configuration—File size, date and time of the download

Secondary Configuration—File size, date and time of the download

Switch Temperature—Either normal or over, for over-temperature

Internal Power Supply—Power supply information. If at full capacity it is displayed in green. If it installed but not operating, it is displayed in red.

External Power Supply—(optional) If present, provides power supply information. If the power supply is operating at full capacity, an OK message displays in green. If it is present, installed, but not operating, the status is displayed in red.

A separate table follows the hardware status that is dedicated to internal cooling fan status.

Figure 92: Hardware Status

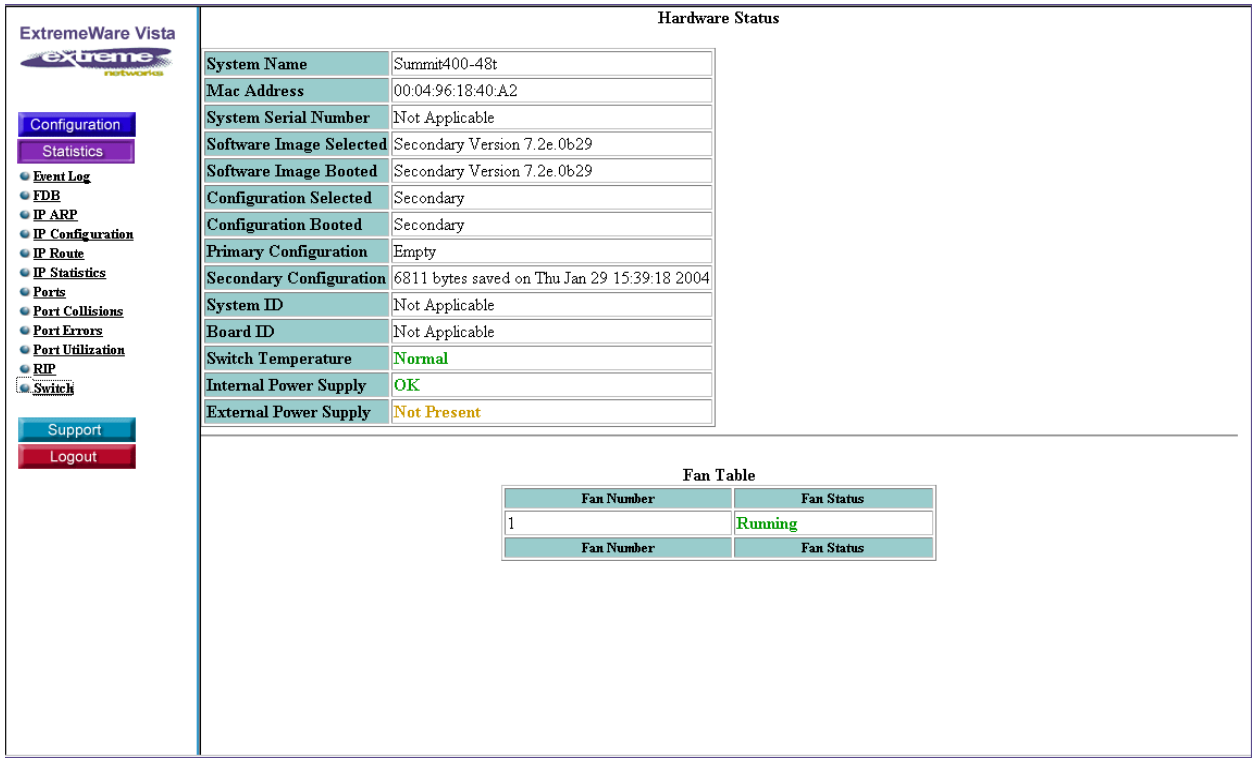

Locating Support Information

ExtremeWare Vista provides a central location to find support information and to download the most current software images. Click **Support** in the task frame to reveal the submenu links:

Help—For links to the most current product manual

TFTP—To upgrade software using a TFTP download

Contact Support—For customer support telephone numbers and URLs

Email Support—To send an email directly to customer support

Help

The Help window provides the URL to the *ExtremeWare 7.2e Installation and User Manual*. See Figure 93 for an example of this window.

Figure 93: Product Manual Link

TFTP Download

You can download the latest software images using Trivial File Transfer Protocol (TFTP) from this window. As shown in Figure 95, you need to provide the following information:

TFTP Server Address—Obtain this address from your Customer Support Representative

Filename—The filename of the software image to download

Container—The location, either primary or secondary, where you want to store the downloaded image

Figure 94: TFTP Download

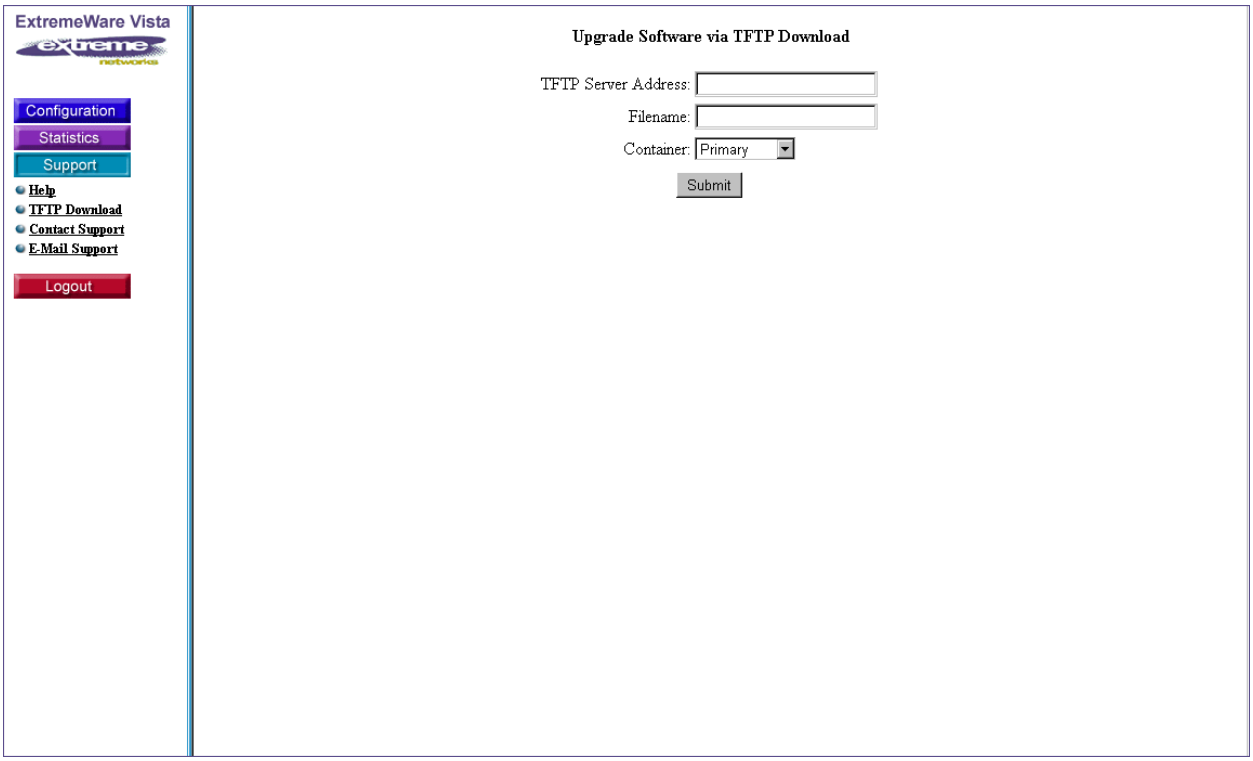

Contact Support

The Contact Support window contains the mailing address, telephone number, fax number, and URL for Customer Support. An example of this window is shown in Figure 95.

Figure 95: Support Address

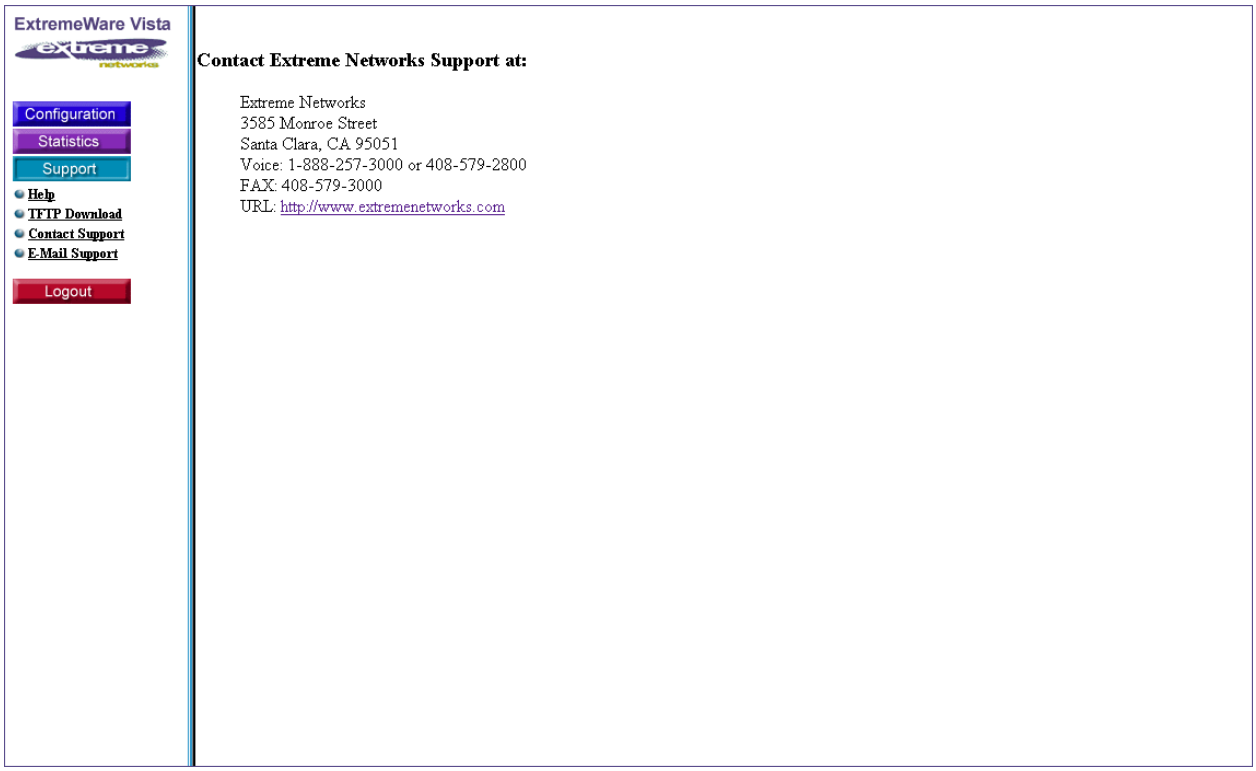

Email Support

When you click the submenu link for Email Support, the browser closes the ExtremeWare Vista page and opens your browser's email window. You can then send an email directly to customer support as shown in Figure 96.

Figure 96: Email Support

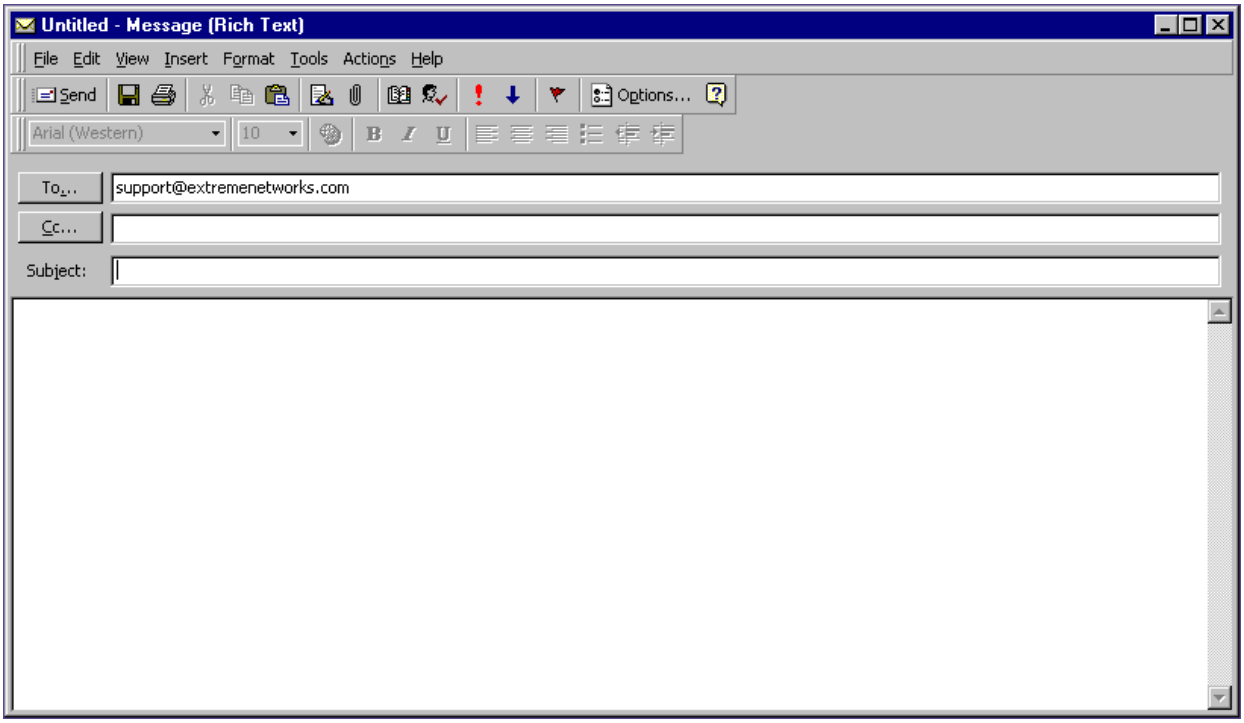

Logging Out of ExtremeWare Vista

When you click the Logout button in the task frame, it causes an immediate exit from ExtremeWare Vista. Be sure you want to exit the application because there is no confirmation screen.

A Technical Specifications

This appendix provides technical specifications for the Summit 400-48 switch. It covers the following topics

Summit 400-48t Switch

The Summit 400-48 has these physical characteristics:

Physical and Environmental	
Dimensions	Height: 1.73 inches (4.40 cm) Width: 17.6 inches (44.1 cm) Depth: 16.4 inches (41.6 cm)
Weight	Weight: 11 lbs (4.98 kg)
Temperature and Humidity	Operating Temperature: 0° to 40° C (32° to 104° F) Storage Temperature: –40° to 70° C (–40° to 158° F) Operating Humidity: 10% to 95% relative humidity, noncondensing Standards: EN60068 to Extreme IEC68 schedule EN 300 019
Power	AC Line Frequency: 50 Hz to 60 Hz Input Voltage Options: 90 VAC to 264 VAC, auto-ranging Current Rating: 4A at 100 VAC; 2A at 240 VAC
Heat Dissipation, Watts/BTU	160 W/0.152 BTU per second

Safety Certifications	
North America	UL 60950 3rd Edition, listed (US Safety) CAN/CSA-C22.2 No. 60950-00 (Canadian Safety)
Europe	Low Voltage Directive (LVD) TUV-R GS Mark by German Notified Body EN60950:2000 (European Safety)
International	CB Scheme IEC60950:2000 with all country deviations (International Safety)
Country Specific	Mexico NOM/NYCE (Product Safety and EMC Approval) Australia/New Zealand AS/NZS 3260 (ACA DoC, Safety of ITE) Argentina S-Mark GOST (Russia)
Laser Safety	
North America	FCC 21 CFR subpart (J) (Safety of Laser Products) CDRH Letter of Approval (US FDA Approval)
Europe	EN60825-2 (European Safety of Lasers)
Electromagnetic Compatibility	
North America	FCC 47 CFR Part 15 Class A (US Emissions) ICES-003 Class A (Canada Emissions)
Europe	89/336/EEC EMC Directive ETSI/EN 300 386:2001 (EU Telecommunications Emissions and Immunity) EN55022:1998 Class A (European Emissions) EN55024:1998 includes IEC/EN 61000-2, 3, 4, 5, 6, 11 (European Immunity) EN 61000-3-2, -3 (Europe Harmonics and Flicker)
International	IEC/CISPR 22:1997 Class A (International Emissions) IEC/CISPR 24:1998 (International Immunity) IEC/EN 61000-4-2 Electrostatic Discharge IEC/EN 61000-4-3 Radiated Immunity IEC/EN 61000-4-4 Transient Bursts IEC/EN 61000-4-5 Surge IEC/EN 61000-4-6 Conducted Immunity IEC/EN 61000-4-11 Power Dips and Interruptions
Country Specific	Japan Class A (VCCI Registration Emissions) Australia/New Zealand AS/NZS 3548 (ACA DoC, Emissions) Korean MIC Mark (MIC Approval, Emissions and Immunity) Mexico NOM/NYCE (Product Safety and EMC Approval) GOST (Russia) Taiwan CNS 13438:1997 Class A (BSMI Approval, Emissions)
Environmental	
Certification Marks	
CE	CE (European Community)
TÜV Rheinland GS geprüfte Sicherheit	TUV/GS (German Notified Body)
TÜV Rheinland Argentina S.A. S	TUV/S (Argentina)

		GOST (Russian Federation)
	ACN 090 029 066	C-Tick (Australian Communication Authority)
		Underwriters Laboratories (USA and Canada)
		MIC (South Korea)
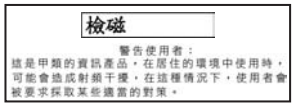		BSMI, Republic of Taiwan
		NOM (Mexican Official Normalization, Electronic Certification and Normalization)

Supported Protocols, MIBs, and Standards

The following is a list of software standards and protocols supported by the Summit 400.

Denial of Service Protection	
RFC 2267 Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing	• CA--96.21.tcp_syn_flooding
RPF (Unicast Reverse Path Forwarding) Control	• CA--96.01.UDP_service_denial
Wire-speed ACLs	• CA--95.01.IP_Spoofing_Attacks_and_Hijacked_ Terminal_Connections
Rate Limiting by ACLs	• CA-2002-03: SNMP vulnerabilities
IP Broadcast Forwarding Control	Host Attacks
ICMP and IP-Option Response Control	• Syndrop
SYN attack protection	• Nestea
Uni-directional Session Control	• Latierra
CERT (http://www.cert.org)	• Newtear
• CA--97.28.Teardrop_Land -Teardrop and “LAND” attack	• Bonk
• IP Options Attack	• Winnuke
• CA--98-13-tcp-denial-of-service	• Raped
• CA--98.01.smurf	• Simping
• CA--96.26.ping	• Sping
	• Ascend
	• Stream

DiffServ - Standards and MIBs
RFC 2474 Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers
RFC 2475 An Architecture for Differentiated Services
RFC 2597 Assured Forwarding PHB Group
RFC 2598 An Expedited Forwarding PHB

Environmental	
EN 300 019-2-1 (2000-09) Storage Class 1.2 - Packaged	ASTM D5276 Drop - Packaged
EN 300 09-2-2 (1999-09) Transportation Class 2.3 - Packaged	ASTM D3332 Shock - Unpackaged
EN 300 019-2-2 (1999-09) Stationary Use at Weather Protected Locations, Class 3.1e - Operational	ASTM D3580 Random Vibration - Unpackaged
EN 300 753 (1997-10) Acoustic Noise - Operational	ASTM D6179 Tilt - Packaged

General Routing and Switching	
RFC 1812 Requirements for IP Version 4 Routers	RFC 826 Ethernet Address Resolution Protocol: Or converting network protocol addresses to 48.bit Ethernet address for transmission on Ethernet hardware
RFC 1519 An Architecture for IP Address Allocation with CIDR	Extreme Standby Router Protocol (ESRP)
RFC 1256 ICMP Router Discovery Messages	IEEE 802.1D-1998 Spanning Tree Protocol
RFC 783 TFTP Protocol (revision 2)	IEEE 802.1W - 2001 Rapid Spanning Tree Protocol
RFC 951 Bootstrap Protocol	IEEE 802.1Q - 1998 Virtual Bridged Local Area Networks
RFC 2131 Dynamic Host Configuration Protocol	Ethernet Automatic Protection Switching (EAPS)-Edge mode, master and member of one ring
RFC 1591 Domain Name System Structure and Delegation	RFC 3619 Ethernet Automatic Protection Switching (EAPS) Version 1
RFC 1122 Requirements for Internet Hosts - Communication Layers	
RFC 768 User Datagram Protocol	
RFC 791 Internet Protocol	
RFC 792 Internet Control Message Protocol	
RFC 793 Transmission Control Protocol	

IP Multicast	
RFC 2362 Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification--two non-passive interfaces	IGMP Snooping with Configurable Router Registration Forwarding
RFC 1112 Host extensions for IP multicasting	Static IGMP Membership
RFC 2236 Internet Group Management Protocol, Version 2	IGMP Filters
	Mtrace, draft-letf-idmr-traceroute-imp-07
	Mrinfo

Management - SNMP & MIBs	
RFC 1157 Simple Network Management Protocol (SNMP)	RFC-1212 Concise MIB definitions
RFC-1215 Convention for defining traps for use with the SNMP	RFC-1213 Management Information Base for Network Management of TCP/IP-based internets: MIB-II
RFC 1573 Evolution of Interface	RFC 1757 Remote Network Monitoring Management Information Base
RFC 1901 Introduction to Community-based SNMPv2	RFC 2021 Remote Network Monitoring Management Information Base Version 2 using SMIv2
RFC 1902 Structure of Management Information for Version 2 of the Simple Network Management Protocol (SNMPv2)	RFC 2613 Remote Network Monitoring MIB Extensions for Switched Networks Version 1.0
RFC 1903 Textual Conventions for Version 2 of the Simple Network Management Protocol (SNMPv2)	RFC 2233 Evolution of the Interfaces Group of MIB-II
RFC 1904 Conformance Statements for Version 2 of the Simple Network Management Protocol (SNMPv2)	RFC 2096 IP Forwarding Table MIB
RFC 1905 Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)	RFC 1724 RIP Version 2 MIB Extension
RFC 1906 Transport Mappings for Version 2 of the Simple Network Management Protocol (SNMPv2)	RFC 1850 OSPF Version 2 Management Information Base
RFC 1907 Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2)	RFC 1155 Structure and identification of management information for TCP/IP-based internets
RFC 1908 Coexistence between Version 1 and Version 2 of the Internet-standard Network Management Framework	RFC 1406 Definitions of Managed Objects for the DS1 and E1 Interface types
RFC 2570 - 2575 SNMPv3, user based security, encryption and authentication	RFC 1407 Definitions of Managed Objects for the DS3/E3 Interface Type
RFC 2576 Coexistence between SNMP Version 1, Version 2 and Version 3	RFC 1493 Definitions of Managed Objects for Bridges
RFC 3410 Introduction and Applicability Statements for Internet-Standard Management Framework	Draft-letf-bridge-rstpmib-03.txt – Definitions of Managed Objects for Bridges with Rapid Spanning Tree Protocol
RFC 3411 An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks	RFC 1354 IPv4 Forwarding Table MIB
RFC 3412 Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)	RFC 2037 Entity MIB
RFC 3413 Simple Network Management Protocol (SNMP) Applications	RFC 1650 Definitions of Managed Objects for the Ethernet-like Interface Types using SMIv2
RFC 3414 User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)	RFC 2665 Definitions of Managed Objects for the Ethernet-like Interface Types
RFC 3415 View-based Access Control Model (VACM) for the Simple Network Management Protocol	RFC 2668 Definitions of Managed Objects for IEEE 802.3 Medium Attachment Units (MAUs)
ExtremeWare vendor MIB (includes ACL, MAC FDB, IP FDB, MAC Address Security, QoS policy and VLAN configuration and statistics, STP and others)	RFC 2787 Definitions of Managed Objects for the Virtual Router Redundancy Protocol
	RFC 2795 Infinite Monkey Protocol Suite
	RFC 2925 Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations
	RFC 1643 Ethernet MIB
	IEEE-802.1x MIB
	Extreme extensions to 802.1x-MIB

Management - Other:	
RFC 1866 Hypertext Markup Language - 2.0	NetFlow version 1 export
RFC 2068 Hypertext Transfer Protocol -- HTTP/1.1	Configuration logging
RFC 854 Telnet Protocol Specification	Multiple Images, Multiple Configs
HTML/ HTTP management	BSD System Logging Protocol (SYSLOG), with Multiple Syslog Servers
Secure Shell 2 (SSH2) client and server	999 Local Messages (criticals stored across reboots)
Secure Copy 2 (SCP2) client and server	RFC 2030 Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI
Telnet client and server	

MPLS - Standards and MIBs	
RFC 2212 Specification of Guaranteed Quality of Service	The Extreme MPLS implementation provides read-only (GET but not SET) support for a subset of the MPLS LSR MIB, as defined in the Internet Draft *draft-ietf-mpls-lsr-mib-07.txt*, and a subset of the MPLS LDP MIB, as defined in the Internet Draft *draft-ietf-mpls-ldp-mib-07.txt*.
RFC 2961 RSVP Overhead Refresh Reduction Extensions	
RFC 3032 MPLS Label Stack Encoding	
RFC 3031 Multiprotocol Label Switching Architecture	
RFC 3036 LDP Specification	
Martini drafts: draft-martini-circuit-encap-mpls-04.txt and draft-martini-l2circuit-trans-mpls-08.txt	
RSVP-TE LSP tunnel draft: draft-ietf-mpls-rsvp-lsp-tunnel-09.txt	
Traffic Engineering Extensions to OSPF: draft-katz-yeung-ospf-traffic-06.txt	

OSPF	
RFC 2328 OSPF Version 2	RFC 1765 OSPF Database Overflow
RFC 1587 The OSPF NSSA Option	RFC 2370 The OSPF Opaque LSA Option

PPP - Standards and MIBs	
RFC 1661 The Point-to-Point Protocol (PPP)	The interface counters in MIB-II (RFC 1213) are supported for PPP. Support for read-only operations (GET operations, but not SET operations) is provided for the following PPP MIBs: • RFC 1471 The Definitions of Managed Objects for the Link Control Protocol of the Point-to-Point Protocol • RFC 1472 The Definitions of Managed Objects for the Security Protocols of the Point-to-Point Protocol • RFC 1474 The Definitions of Managed Objects for the Bridge Network Control Protocol of the Point-to-Point Protocol • RFC 1473 The Definitions of Managed Objects for the IP Network Control Protocol of the Point-to-Point Protocol
RFC 1662 PPP in HDLC-like Framing	
RFC 2615 PPP over SONET/SDH	
RFC 1334 PPP Authentication Protocols	
RFC 1994 PPP Challenge Handshake Authentication Protocol (CHAP)	
RFC 1332 The PPP Internet Protocol Control Protocol (IPCP)	
RFC 2878 PPP Bridging Control Protocol (BCP)	
RFC 1191 Path MTU Discovery	
RFC 3032 MPLS Label Stack Encoding	

Quality of Service	
IEEE 802.1D -1998 (802.1p) Packet Priority	RFC 2475 An Architecture for Differentiated Service
RFC 2474 Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers	Layer 1-4, layer 7 (user name) Policy-Based Mapping
RFC 2598 An Expedited Forwarding PHB	Policy-Based Mapping/Overwriting of DiffServ code points, .1p priority
RFC 2597 Assured Forwarding PHB Group	DLCS (Dynamic Link Context System, WINS snooping) for integration with EPICenter Policy Manager
Bi-directional Rate Shaping	

RIP	
RFC 1058 Routing Information Protocol	RFC 2453 RIP Version 2

Security	
Routing protocol authentication (see above)	Multiple supplicants for Network Login (web-based and 802.1x modes)
Secure Shell (SSHv2) & Secure Copy (SCPv2) with encryption/authentication	RADIUS Per-command Authentication
SNMPv3 user based security, with encryption/authentication	Access Profiles on All Routing Protocols
RFC 1492 An Access Control Protocol, Sometimes Called TACACS	Access Profiles on All Management Methods
RFC 2138 Remote Authentication Dial In User Service (RADIUS)	Network Login (including DHCP / RADIUS integration)
RFC 2139 RADIUS Accounting	MAC Address Security / Lockdown
IEEE 802.1x Port Based Network Access Control	Network Address Translation (NAT)
	Layer 2/3/4/7 Access Control Lists (ACLs)

VLANs	
IEEE 802.1Q VLAN Tagging	Multiple STP domains per VLAN
IEEE 802.3ad Static ConfigPort-based VLANs	RFC 3069 VLAN Aggregation for Efficient IP Address Allocation
IEEE 802.1v VLAN classification by Protocol and Port	VLAN Translation
Port-based VLANs	RFC 2674 Definitions of Managed Objects for Bridges with Traffic Classes, Multicast Filtering, and Virtual LAN Extensions
MAC-based VLANs	
Virtual MANs	

B Software Upgrade and Boot Options

This appendix describes the following topics:

Downloading a New Image

The image file contains the executable code that runs on the switch. It comes preinstalled from the factory. As new versions of the image are released, you should upgrade the software running on your system.

The image is upgraded by using a download procedure from either a Trivial File Transfer Protocol (TFTP) server on the network or from a PC connected to the serial port using the XMODEM protocol. Downloading a new image involves the following steps:

- Load the new image onto a TFTP server on your network (if you will be using TFTP).
- Load the new image onto a PC (if you will be using XMODEM).
- Download the new image to the switch using the following command:

  ```
  download image [<hostname> | <ipaddress>] <filename>  {primary | secondary}
  ```

  where the following is true:

  `hostname`—Is the hostname of the TFTP server. (You must enable DNS to use this option.)

  `ipaddress`—Is the IP address of the TFTP server.

  `filename`—Is the filename of the new image.

  `primary`—Indicates the primary image.

  `secondary`—Indicates the secondary image.

Selecting a Primary or a Secondary Image

The switch can store up to two images: a primary and a secondary. When you download a new image, you must select into which image space (primary or secondary) the new image should be placed. If not

indicated, the next selected boot-up image space is used. This is the primary image space by default, but it can be changed with the following command:

```
use image [primary | secondary]
```

Understanding the Image Version String

The image version string contains build information for each version of ExtremeWare. You can use either the `show version` or `show switch` command to display the ExtremeWare version running on your switch.

Depending on the CLI command, the output is structured as follows:

- `show version`

  Version <major>.<sub_major>.<minor> (Build<build>) {[branch | beta | tech | patch]{<image_version>}.<image_description>-r<branch_revision>}

- `show switch`

  <major>.<sub_major>.<minor>b<build>{[branch | beta | tech | patch]{<image_version>}.<image_description>-r<branch_revision>}

Table 42 describes the image version fields.

Table 42: Image version fields

Field	Description
major	Specifies the ExtremeWare Major version number.
sub_major	Specifies the ExtremeWare Sub-major version number.
minor	Specifies the ExtremeWare Minor version number.
build	Specifies the ExtremeWare build number. This value is reset to zero for each new Major and Minor release.
image_version	Identifies the Technology Release or Beta image version. The image version number is zero for all but Technology Releases and Beta releases.
image_description	Identifies a specific Patch, Beta Release, Technology Release, or Development Branch Release.
branch_revision	Indicates an incremental build on a specific branch. The branch revision number is zero for General Availability and Sustaining releases.

Table 43 displays sample `show version` and `show switch` output for various ExtremeWare versions.

Table 43: Sample show output

Release Type	Show Version Command	Show Switch Command
Major	`Version 7.0.0 (Build 61)`	`7.0.0b61`
Minor	`Version 7.0.1 (Build 4)`	`7.0.1b4`
Sustaining	`Version 7.0.0 (Build 68)`	`7.0.0b68`
Patch	`Version 7.0.0 (Build 61) patch.030131-01-r1`	`7.0.0b61 patch.030131-01-r1`
Technology	`Version 7.0.0 (Build 68) tech2.ipv6-r4`	`7.0.0b68 tech2.ipv6-r4`

Table 43: Sample show output

Release Type	Show Version Command	Show Switch Command
Beta	`Version 7.0.1 (Build 3) beta1.triumph-r4`	`7.0.1b3 beta1.triumph-r4`
Development Branch	`Version 7.0.0 (Build 67) branch.triumph-r5`	`7.0.0b67 branch.triumph-r5`

Software Signatures

Each ExtremeWare image contains a unique signature. The BootROM checks for signature compatibility and denies an incompatible software upgrade. In addition, the software checks both the installed BootROM and software and also denies an incompatible upgrade.

Rebooting the Switch

To reboot the switch, use the following command:

```
reboot {time <date> <time> | cancel}
```

where `date` is the date and `time` is the time (using a 24-hour clock format) when the switch will be rebooted. The values use the following format:

```
mm/dd/yyyy hh:mm:ss
```

If you do not specify a reboot time, the reboot occurs immediately following the command, and any previously schedule reboots are cancelled. To cancel a previously scheduled reboot, use the `cancel` option.

Saving Configuration Changes

The configuration is the customized set of parameters that you have selected to run on the switch. As you make configuration changes, the new settings are stored in run-time memory. Settings that are stored in run-time memory are not retained by the switch when the switch is rebooted. To retain the settings, and have them load when you reboot the switch, you must save the configuration to nonvolatile storage.

The switch can store two different configurations: a primary and a secondary. When you save configuration changes, you can select to which configuration you want the changes saved. If you do not specify, the changes are saved to the configuration area currently in use.

To save the configuration, use the following command:

```
save configuration {primary | secondary}
```

To use the configuration, use the following command:

```
use configuration [primary | secondary]
```

The configuration takes effect on the next reboot.

NOTE

If the switch is rebooted while in the middle of a configuration save, the switch boots to factory default settings. The configuration that is not in the process of being saved is unaffected.

Returning to Factory Defaults

To return the switch configuration to factory defaults, use the following command:

```
unconfigure switch
```

This command resets the entire configuration, with the exception of user accounts and passwords that have been configured, and the date and time.

To erase the currently selected configuration image and reset all switch parameters, use the following command:

```
unconfigure switch {all}
```

Using TFTP to Upload the Configuration

You can upload the current configuration to a TFTP server on your network. The uploaded ASCII file retains the command-line interface (CLI) format. This allows you to:

- Modify the configuration using a text editor, and later download a copy of the file to the same switch, or to one or more different switches.
- Send a copy of the configuration file to the Extreme Networks Technical Support department for problem-solving purposes.
- Automatically upload the configuration file every day, so that the TFTP server can archive the configuration on a daily basis. Because the filename is not changed, the configured file stored in the TFTP server is overwritten every day.

To upload the configuration, use the following command:

```
upload configuration [<ip address> | <hostname>] <filename> {every <time>}
```

where the following is true:

- `ipaddress`—Is the IP address of the TFTP server.
- `hostname`—Is the hostname of the TFTP server. (You must enable DNS to use this option.)
- `filename`—Is the name of the ASCII file. The filename can be up to 255 characters long, and cannot include any spaces, commas, quotation marks, or special characters.
- `every <time>`—Specifies the time of day you want the configuration automatically uploaded on a daily basis. If not specified, the current configuration is immediately uploaded to the TFTP server.

To cancel a previously scheduled configuration upload, use the following command:

```
upload configuration cancel
```

Using TFTP to Download the Configuration

You can download ASCII files that contain CLI commands to the switch to modify the switch configuration. Three types of configuration scenarios that can be downloaded:

- Complete configuration
- Incremental configuration
- Scheduled incremental configuration

If you load a configuration from a different model, you can safely write the correct configuration over the unsupported configuration.

Downloading a Complete Configuration

Downloading a complete configuration replicates or restores the entire configuration to the switch. You typically use this type of download in conjunction with the `upload configuration` command, which generates a complete switch configuration in an ASCII format. As part of the complete configuration download, the switch is automatically rebooted.

To download a complete configuration, use the `download configuration` command using the following syntax:

```
download configuration [<ip address> | <hostname>] <filename>
```

After the ASCII configuration is downloaded by way of TFTP, you are prompted to reboot the switch. The downloaded configuration file is stored in current switch memory during the rebooting process, and is not retained if the switch has a power failure.

When the switch completes booting, it treats the downloaded configuration file as a script of CLI commands, and automatically executes the commands. If your CLI connection is through a Telnet connection (and not the console port), your connection is terminated when the switch reboots, but the command executes normally.

Downloading an Incremental Configuration

A partial or incremental change to the switch configuration may be accomplished by downloaded ASCII files that contain CLI commands. These commands are interpreted as a script of CLI commands, and take effect at the time of the download, without requiring a reboot of the switch.

To download an incremental configuration, use the following command:

```
download configuration [<ip address> | <hostname>] <filename> {incremental}
```

Do not download an incremental configuration when you have time-critical applications running. When you download an incremental configuration, the switch immediately processes the changes, which can affect the processing of other tasks. We recommend that you either download small incremental configurations, or schedule downloads during maintenance windows.

Scheduled Incremental Configuration Download

You can schedule the switch to download a partial or incremental configuration on a regular basis. You could use this feature to update the configuration of the switch regularly from a centrally administered TFTP server. As part of the scheduled incremental download, you can optionally configure a backup TFTP server.

To configure the primary and/or secondary TFTP server and filename, use the following command:

```
configure download server [primary | secondary] [<ip address> | <hostname>] <filename>
```

To enable scheduled incremental downloads, use the following command:

```
download configuration every <time>
```

To display scheduled download information, use the following command:

```
show switch
```

To cancel scheduled incremental downloads, use the following command:

```
download configuration cancel
```

Remember to Save

Regardless of which download option is used, configurations are downloaded into switch runtime memory, only. The configuration is saved only when the save command is issued, or if the configuration file, itself, contains the save command.

If the configuration currently running in the switch does not match the configuration that the switch used when it originally booted, an asterisk (*) appears before the command line prompt when using the CLI.

Upgrading and Accessing BootROM

The BootROM of the switch initializes certain important switch variables during the boot process. If necessary, BootROM can be upgraded, after the switch has booted, using TFTP. In the event the switch does not boot properly, some boot option functions can be accessed through a special BootROM menu.

Upgrading BootROM

Upgrading BootROM is done using TFTP (from the CLI), after the switch has booted. Upgrade the BootROM only when asked to do so by an Extreme Networks technical representative. To upgrade the BootROM, use the following command:

```
download bootrom [<ip address> | <hostname>] <filename>
```

Accessing the BootROM Menu

Interaction with the BootROM menu is only required under special circumstances, and should be done only under the direction of Extreme Networks Customer Support. The necessity of using these functions implies a non-standard problem which requires the assistance of Extreme Networks Customer Support.

To access the BootROM menu, follow these steps:

1. Attach a serial cable to the console port of the switch.
2. Attach the other end of the serial cable to a properly configured terminal or terminal emulator, power cycle the switch while depressing the spacebar on the keyboard of the terminal.

   As soon as you see the `BootROM->` prompt, release the spacebar. You can see a simple help menu by pressing `h`. Options in the menu include

   - Selecting the image to boot from
   - Booting to factory default configuration
   - Performing a serial download of an image

For example, to change the image that the switch boots from in flash memory, press `1` for the image stored in primary or `2` for the image stored in secondary. Then, press the `f` key to boot from newly selected on-board flash memory.

To boot to factory default configuration, press the `d` key for default and the `f` key to boot from the configured on-board flash.

To perform a serial download, you can optionally change the baud rate to 115200 using the `b` command. Then press the `s` key to prepare the switch for an image to be sent from your terminal using the 1K XMODEM protocol. (You can use use a Windows Hyperterminal program to accomplish this step.) After the transfer is complete, the switch restores the console port to 9600 bps and begins the boot process.

C Troubleshooting

If you encounter problems when using the switch, this appendix may be helpful. If you have a problem not listed here or in the release notes, contact your local technical support representative.

LEDs

Power LED does not light:

Check that the power cable is firmly connected to the device and to the supply outlet.

On powering-up, the MGMT LED lights yellow:

The device has failed its Power On Self Test (POST) and you should contact your supplier for advice.

A link is connected, but the Status LED does not light:

Check that:

- All connections are secure.
- Cables are free from damage.
- The devices at both ends of the link are powered-up.
- Both ends of the Gigabit link are set to the same autonegotiation state.

  The Gigabit link must be enabled or disabled on both sides. If the two sides are different, typically the side with autonegotiation disabled will have the link LED lit, and the side with autonegotiation enabled will not be lit. The default configuration for a Gigabit port is autonegotiation enabled. This can be verified by entering the following command:

  ```
  show ports {mgmt | <portlist>} configuration
  ```

Switch does not power up:

All products manufactured by Extreme Networks use digital power supplies with surge protection. In the event of a power surge, the protection circuits shut down the power supply. To reset, unplug the switch for 1 minute, plug it back in, and attempt to power up the switch.

If this does not work, try using a different power source (different power strip / outlet) and power cord.

Cable Diagnostics

If you are having a problem establishing a link, you might have a faulty Ethernet cable. An Ethernet cable is composed of four pairs of unshielded twisted-pair (UTP). Of those four pairs, two are required to create the link. In addition to physically inspecting the cable, you can run a CLI command to test the cable. Use the following commands to test an Ethernet cable and to display the output of the test:

```
run diagnostics cable ports <portlist>
```

```
show ports <portlist> cable diagnostics
```

The `run diagnostics cable ports` command prompts you when the diagnostics are complete to enter the `show ports cable diagnostics` command.

By reviewing the output of the `show` command you can determine:

- The length of the cable
- Whether there is a successful termination, or whether there is an open or short

For example, the following command set tests the Ethernet cable inserted into port 1. The four copper pairs do not all have the same length, which might indicate a kink in the cable, or a open connection:

```
Summit400-48t:27 # run diagnostics cable ports 1

Cable Diagnostics has completed, to view results enter

show port <port list> cable diagnostics

Summit400-48t:28 # show port 1 cable diagnostics

Port       Pair    Length                          Status

   1    Pair A    3 meters          Terminated
        Pair B    2 meters          Terminated
        Pair C    1 meters          Open or Short
        Pair D    1 meters          Open or Short
```

The next example shows none of the twisted pairs terminate successfully at port 1, which could indicate that the cable is not inserted into the port:

```
Summit400-48t:29 # run diagnostics cable ports 1

Cable Diagnostics has completed, to view results enter

show port <port list> cable diagnostics

Summit400-48t:30 # show port 1 cable diagnostics

Port       Pair    Length                          Status

   1    Pair A    0 meters          Open or Short
        Pair B    0 meters          Open or Short
        Pair C    0 meters          Open or Short
        Pair D    0 meters          Open or Short
```

Using the Command-Line Interface

The initial welcome prompt does not display:

Check that your terminal or terminal emulator is correctly configured.

For console port access, you may need to press [Return] several times before the welcome prompt appears.

Check the settings on your terminal or terminal emulator. The settings are 9600 baud, 8 data bits, 1 stop bit, no parity, XON/OFF flow control enabled.

The SNMP Network Manager cannot access the device:

Check that the device IP address, subnet mask, and default router are correctly configured, and that the device has been reset.

Check that the device IP address is correctly recorded by the SNMP Network Manager (refer to the user documentation for the Network Manager).

Check that the community strings configured for the system and Network Manager are the same.

Check that the SNMPv3 USM, Auth, and VACM configured fore the system and Network Manager are the same.

Check that SNMP access was not disabled for the system.

The Telnet workstation cannot access the device:

Check that the device IP address, subnet mask and default router are correctly configured, and that the device has been reset. Ensure that you enter the IP address of the switch correctly when invoking the Telnet facility. Check that Telnet access was not disabled for the switch. If you attempt to log in and the maximum number of Telnet sessions are being used, you should receive an error message indicating so.

Traps are not received by the SNMP Network Manager:

Check that the SNMP Network Manager's IP address and community string are correctly configured, and that the IP address of the Trap Receiver is configured properly on the system.

The SNMP Network Manager or Telnet workstation can no longer access the device:

Check that Telnet access or SNMP access is enabled.

Check that the port through which you are trying to access the device has not been disabled. If it is enabled, check the connections and network cabling at the port.

Check that the port through which you are trying to access the device is in a correctly configured VLAN.

Try accessing the device through a different port. If you can now access the device, a problem with the original port is indicated. Re-examine the connections and cabling.

A network problem may be preventing you accessing the device over the network. Try accessing the device through the console port.

Check that the community strings configured for the device and the Network Manager are the same.

Check that SNMP access was not disabled for the system.

Permanent entries remain in the FDB:

If you have made a permanent entry in the FDB (which requires you to specify the VLAN to which it belongs and then delete the VLAN), the FDB entry will remain. Though causing no harm, you must manually delete the entry from the FDB if you want to remove it.

Default and Static Routes:

If you have defined static or default routes, those routes will remain in the configuration independent of whether the VLAN and VLAN IP address that used them remains. You should manually delete the routes if no VLAN IP address is capable of using them.

You forget your password and cannot log in:

If you are not an administrator, another user having administrator access level can log in, delete your user name, and create a new user name for you, with a new password.

Alternatively, another user having administrator access level can log in and initialize the device. This will return all configuration information (including passwords) to the initial values.

In the case where no one knows a password for an administrator level user, contact your supplier.

Port Configuration

No link light on 10/100/1000 Base port:

If patching from a hub or switch to another hub or switch, ensure that you are using a CAT5 cross-over cable. This is a CAT5 cable that has pins 1&2 on one end connected to pins 3&6 on the other end. Also try running the cable diagnostics, as described in “Cable Diagnostics” on page 322.

Excessive RX CRC errors:

When a device that has auto-negotiation disabled is connected to a Extreme switch that has auto-negotiation enabled, the Extreme switch links at the correct speed, but in half duplex mode. The Extreme switch 10/100/1000 physical interface uses a method called *parallel detection* to bring up the link. Because the other network device is not participating in auto-negotiation (and does not advertise its capabilities), parallel detection on the Extreme switch is only able to sense 10Mbps versus 100Mbps speed, and not the duplex mode. Therefore, the switch establishes the link in half duplex mode using the correct speed.

The only way to establish a full duplex link is to either force it at both sides, or run auto-negotiation on both sides (using full duplex as an advertised capability, which is the default setting on the Extreme switch).

NOTE

A mismatch of duplex mode between the Extreme switch and another network device will cause poor network performance. Viewing statistics using the `show ports rxerrors` *command on the Extreme switch may display a constant increment of CRC errors. This is characteristic of a duplex mismatch between devices. This is NOT a problem with the Extreme switch.*

Always verify that the Extreme switch and the network device match in configuration for speed and duplex.

No link light on Gigabit fiber port:

Check to ensure that the transmit fiber goes to the receive fiber side of the other device, and vice-versa. All gigabit fiber cables are of the cross-over type.

The Extreme switch has auto-negotiation set to on by default for gigabit ports. These ports need to be set to auto off (using the command `configure port <port #> auto off`) if you are connecting it to devices that do not support auto-negotiation.

Ensure that you are using multi-mode fiber (MMF) when using a 1000BASE-SX GBIC, and single mode fiber (SMF) when using a 1000BASE-LX GBIC. 1000BASE-SX does not work with SMF. 1000BASE-LX works with MMF, but requires the use of a mode conditioning patchcord (MCP).

VLANs

You cannot add a port to a VLAN:

If you attempt to add a port to a VLAN and get an error message similar to

```
localhost:7 # configure vlan marketing add port 1:1,1:2
ERROR: Protocol conflict on port 1:5
```

you already have a VLAN using untagged traffic on a port. Only one VLAN using untagged traffic can be configured on a single physical port.

VLAN configuration can be verified by using the following command:

```
show vlan {<vlan name> | detail | stats {vlan} <vlan name>}
```

The solution for this error is to remove ports 1 and 2 from the VLAN currently using untagged traffic on those ports. If this were the “default” VLAN, the command would be

```
localhost # configure vlan default del port 1,2
```

which should now allow you to re-enter the previous command without error as follows:

```
localhost # configure vlan add port 1,2
```

VLAN names:

There are restrictions on VLAN names. They cannot contain whitespaces and cannot start with a numeric value unless you use quotation marks around the name. If a name contains whitespaces, starts with a number, or contains non-alphabetical characters, you must use quotation marks whenever referring to the VLAN name.

802.1Q links do not work correctly:

Remember that VLAN names are only locally significant through the command-line interface. For two switches to communicate across a 802.1Q link, the VLAN ID for the VLAN on one switch should have a corresponding VLAN ID for the VLAN on the other switch.

If you are connecting to a third-party device and have checked that the VLAN IDs are the same, the Ethertype field used to identify packets as 802.1Q packets may differ between the devices. The default value used by the switch is 8100.

VLANs, IP Addresses and default routes:

The system can have an IP address for each configured VLAN. It is necessary to have an IP address associated with a VLAN if you intend to manage (Telnet, SNMP, ping) through that VLAN or route IP traffic. You can also configure multiple default routes for the system. The system first tries the default route with the lowest cost metric.

STP

You have connected an endstation directly to the switch and the endstation fails to boot correctly:

The switch has STP enabled, and the endstation is booting before the STP initialization process is complete. Specify that STP has been disabled for that VLAN, or turn off STP for the switch ports of the endstation and devices to which it is attempting to connect, and then reboot the endstation.

The switch keeps aging out endstation entries in the switch Forwarding Database (FDB):

Reduce the number of topology changes by disabling STP on those systems that do not use redundant paths.

Specify that the endstation entries are static or permanent.

Debug Tracing/Debug Mode

In ExtremeWare version 7.1.0, the Event Management System (EMS) facility was added to ExtremeWare. EMS provides a standardized way to filter and store messages generated by the switch. Many of the systems in ExtremeWare are moving into EMS. As a system is converted to EMS, the corresponding debug trace command associated with that system is removed. With EMS, you must enable debug mode to display debug information. To enable or disable debug mode for EMS, use the following commands:

```
enable log debug-mode
disable log debug-mode
```

Once debug mode is enabled, you can configure EMS to capture specific debug information from the switch. Details of EMS can be found in Chapter 8, “Status Monitoring and Statistics” on page 119.

For the systems not yet converted to EMS, ExtremeWare includes a debug tracing facility for the switch. The `show debug-trace` command can be applied to one or all VLANs, as follows:

```
show debug-trace {vlan <vlan name>}
```

The debug commands should only be used under the guidance of Extreme Networks technical personnel.

To reset all debug-tracing to the default level, use the following command:

```
clear debug-trace
```

To change the debug tracing facility for a certain system to a specified debug level, use the following command:

```
configure debug-trace <system> <level> vlan <vlan name>
```

Some of the debug trace systems commands can be applied to a particular VLAN, some apply to the switch as a whole, so the `vlan` option is not available with all systems.

To display the debug tracing configuration, use the following command:

```
show debug-trace <system> vlan <vlan name>
```

Again, the `vlan` option is not available with every system.

TOP Command

The `top` command is a utility that indicates CPU utilization by process.

System Memory Dump

You can download the entire contents of memory through the Ethernet management port. This is used only for troubleshooting, and should not be used without assistance from Extreme Networks technical support.

To specify the IP address to which to transfer a dump if the `system-dump` option is specified in the configuration, use the following command:

```
configure system-dump server <ip address>
```

This address is also used if no address is provided in the command. The default is 0 or “no IP”.

To set an optional timeout for the dump transfer, use the following command:

```
configure system-dump timeout <seconds>
```

The default is 0. The minimum non-zero value is 120 seconds. The minimum recommended value is 480 seconds.

To display the system-dump server IP and dump-timeout, use the following command:

```
show system-dump
```

To specify a memory dump if a task generates a software exception, use the following command:

```
configure sys-recovery-level [none | [all | critical] [ reboot | shutdown |
system-dump [maintenance-mode | reboot | shutdown]]]
```

The four options specify the action taken when the dump transfer is complete. The actions occur whether or not the dump was successful. The `maintenance-mode` option leaves the switch in whatever state it was in before the dump.

System Odometer

Each field replaceable component contains a system odometer counter in EEPROM. You can use the `show switch` command to see how long an individual component has been in service since it was manufactured.

Reboot Loop Protection

If the system reboots due to a failure that remains after the reboot, it reboots when it detects the failure again. To protect against continuous reboot loops, you can configure reboot loop protection using the following command:

```
configure reboot-loop-protection threshold <time-interval> <count>
```

If the switch reboots the specified number of times within the specified time interval, it stops rebooting and comes up in minimal mode. If you reboot the switch manually or `run diagnostics` commands, the time interval and count are both reset to 0.

Minimal Mode

In minimal mode, only the CPU, NVRAM, management port, and minimal tasks are active. The following commands are supported in minimal mode:

- reboot
- unconfigure switch all
- unconfigure switch
- use image
- use configuration
- download bootrom
- download image
- download configuration
- configure iparp
- configure vlan ipaddress
- configure iproute add default
- configure diagnostics
- show iproute
- show iparp
- show vlan
- show version
- show log
- ping
- clear log
- clear log diag-status

Contacting Extreme Technical Support

If you have a network issue that you are unable to resolve, contact Extreme Networks technical support. Extreme Networks maintains several Technical Assistance Centers (TACs) around the world to answer networking questions and resolve network problems. You can contact technical support by phone at:

- (800) 998-2408
- (408) 579-2826

or by email at:

- support@extremenetworks.com

You can also visit the support website at:

http://www.extremenetworks.com/services/resources/

to download software updates (requires a service contract) and documentation (including a.pdf version of this manual).

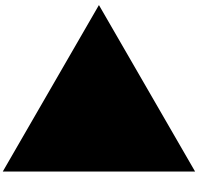

Index of Commands

D

E

H

L

M

N

P

Q

R

S

T

U

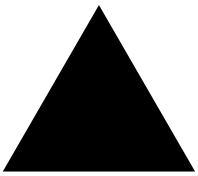

Index

C

D

E

F

G

H

I

S

T

U

V

W

X

Z