
Summit24e3 Switch Installation and User Guide

Software Version 2.0

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

Published: August 2002
Part number: 100102-00 Rev. 02

©2002 Extreme Networks, Inc. All rights reserved. Extreme Networks and BlackDiamond are registered trademarks of Extreme Networks, Inc. in the United States and certain other jurisdictions. ExtremeWare, ExtremeWare Vista, ExtremeWorks, ExtremeAssist, ExtremeAssist1, ExtremeAssist2, PartnerAssist, Extreme Standby Router Protocol, ESRP, SmartTraps, Alpine, Summit, Summit1, Summit4, Summit4/FX, Summit7i, Summit24, Summit48, Summit Virtual Chassis, SummitLink, SummitGbX, SummitRPS and the Extreme Networks logo are trademarks of Extreme Networks, Inc., which may be registered or pending registration in certain jurisdictions. The Extreme Turbodriven logo is a service mark of Extreme Networks, which may be registered or pending registration in certain jurisdictions. Specifications are subject to change without notice.

NetWare and Novell are registered trademarks of Novell, Inc. Merit is a registered trademark of Merit Network, Inc. Solaris is a trademark of Sun Microsystems, Inc. F5, BIG/ip, and 3DNS are registered trademarks of F5 Networks, Inc. see/IT is a trademark of F5 Networks, Inc.

 "Data Fellows", the triangle symbol, and Data Fellows product names and symbols/logos are trademarks of Data Fellows.

 F-Secure SSH is a registered trademark of Data Fellows.


All other registered trademarks, trademarks and service marks are property of their respective owners.

Authors: Megan Mahar, Hugh Bussell

Editor:

Production:

Special Thanks: Andy Jeyanandan, Victor Tseng



Contents

	Preface	
	Introduction	17
	Conventions	17
	Related Publications	18
Chapter 1	Summit24e3 Switch Overview	
	Summary of Features	19
	Summit24e3 Switch Front View	20
	Console Port	20
	Reset Button	20
	Summit24e3 Switch Rear View	21
	Power Socket	21
	Serial Number	21
	MAC Address	21
	Summit24e3 Switch LEDs	21
	Port Connections	22
	Full-Duplex	22
	Mini-GBIC Type and Hardware/Software Support	22
	Mini-GBIC Type and Specifications	22
	Safety Information	23
Chapter 2	Switch Installation	
	Determining the Switch Location	25
	Following Safety Information	26
	Installing the Switch	26
	Rack Mounting	26
	Free-Standing	27
	Stacking the Switch and Other Devices	27

Connecting Equipment to the Console Port	27
Powering On the Switch	28
Checking the Installation	28
Logging In for the First Time	29
Installing or Replacing a Mini-Gigabit Interface Connector (Mini-GBIC)	30
Safety Information	30
Preparing to Install or Replace a Mini-GBIC	30
Installing and Removing a Mini-GBIC	31
 Chapter 3 ExtremeWare Overview	
Summary of Features	33
Virtual LANs (VLANs)	34
Spanning Tree Protocol	34
Quality of Service	35
Unicast Routing	35
Load Sharing	35
ESRP-Aware Switches	35
Software Licensing	36
Feature Licensing	36
Security Licensing	37
Obtaining a Security License	37
Security Features Under License Control	37
Software Factory Defaults	38
 Chapter 4 Accessing the Switch	
Understanding the Command Syntax	39
Syntax Helper	40
Command Shortcuts	40
Summit24e3 Switch Numerical Ranges	40
Names	41
Symbols	41
Line-Editing Keys	41
Command History	42
Common Commands	42
Configuring Management Access	44
User Account	44
Administrator Account	45
Default Accounts	45
Creating a Management Account	46
Domain Name Service Client Services	47

	Checking Basic Connectivity	47
	Ping	48
	Traceroute	48
Chapter 5	Managing the Switch	
	Overview	49
	Using the Console Interface	50
	Using Telnet	50
	Connecting to Another Host Using Telnet	50
	Configuring Switch IP Parameters	50
	Disconnecting a Telnet Session	52
	Controlling Telnet Access	52
	Using Secure Shell 2 (SSH2)	53
	Enabling SSH2	53
	Using SNMP	54
	Accessing Switch Agents	54
	Supported MIBs	54
	Configuring SNMP Settings	54
	Displaying SNMP Settings	55
	Authenticating Users	56
	RADIUS Client	56
	Configuring TACACS+	60
	Using Network Login	62
	Using Network Login in Campus Mode	62
	Using Network Login in ISP Mode	65
	DHCP Server on the Switch	66
	Network Login Configuration Commands	66
	Displaying Network Login Settings	66
	Disabling Network Login	67
	Using the Simple Network Time Protocol	67
	Configuring and Using SNTP	67
	SNTP Configuration Commands	69
	SNTP Example	70
Chapter 6	Configuring Ports on a Switch	
	Enabling and Disabling Switch Ports	71
	Configuring Switch Port Speed and Duplex Setting	71
	Switch Port Commands	72
	Load Sharing on the Switch	73
	Load-Sharing Algorithms	73
	Configuring Switch Load Sharing	74
	Load-Sharing Example	75

Verifying the Load-Sharing Configuration	75
Switch Port-Mirroring	75
Port-Mirroring Commands	76
Port-Mirroring Example	76
Extreme Discovery Protocol	76
EDP Commands	77
 Chapter 7 Virtual LANs (VLANs)	
Overview of Virtual LANs	79
Benefits	79
Types of VLANs	80
Port-Based VLANs	80
Tagged VLANs	82
VLAN Names	84
Default VLAN	84
Renaming a VLAN	85
Configuring VLANs on the Switch	85
VLAN Configuration Commands	85
VLAN Configuration Examples	86
Displaying VLAN Settings	86
 Chapter 8 Forwarding Database (FDB)	
Overview of the FDB	87
FDB Contents	87
FDB Entry Types	87
How FDB Entries Get Added	88
Associating a QoS Profile with an FDB Entry	88
Configuring FDB Entries	88
FDB Configuration Examples	90
Displaying FDB Entries	90
 Chapter 9 Access Policies	
Overview of Access Policies	91
Access Control Lists	91
Rate Limits	91
Routing Access Policies	92
Using Access Control Lists	92
Access Masks	92
Access Lists	92
Rate Limits	93
How Access Control Lists Work	94

	Access Mask Precedence Numbers	94
	Specifying a Default Rule	94
	The permit-established Keyword	95
	Adding Access Mask, Access List, and Rate Limit Entries	95
	Deleting Access Mask, Access List, and Rate Limit Entries	96
	Verifying Access Control List Configurations	96
	Access Control List Commands	96
	Access Control List Examples	100
	Using Routing Access Policies	104
	Creating an Access Profile	104
	Configuring an Access Profile Mode	104
	Adding an Access Profile Entry	104
	Deleting an Access Profile Entry	105
	Applying Access Profiles	105
	Routing Access Policies for RIP	105
	Routing Access Policies for OSPF	107
	Making Changes to a Routing Access Policy	108
	Removing a Routing Access Policy	108
	Routing Access Policy Commands	109
Chapter 10	Network Address Translation (NAT)	
	Overview	111
	Internet IP Addressing	112
	Configuring VLANs for NAT	112
	NAT Modes	113
	Configuring NAT	114
	Configuring NAT Rules	114
	Creating NAT Rules	114
	Creating Static and Dynamic NAT Rules	115
	Creating Portmap NAT Rules	115
	Creating Auto-Constrain NAT Rules	115
	Advanced Rule Matching	116
	Configuring Timeouts	116
	Displaying NAT Settings	117
	Disabling NAT	117
Chapter 11	Ethernet Automatic Protection Switching	
	Overview of the EAPS Protocol	119
	Fault Detection and Recovery	121
	Restoration Operations	123

Commands for Configuring and Monitoring EAPS	123
Creating and Deleting an EAPS Domain	124
Defining the EAPS Mode of the Switch	125
Configuring EAPS Polling Timers	125
Configuring the Primary and Secondary Ports	126
Configuring the EAPS Control VLAN	126
Configuring the EAPS Protected VLANs	127
Enabling and Disabling an EAPS Domain	127
Enabling and Disabling EAPS	128
Unconfiguring an EAPS Ring Port	128
Displaying EAPS Status Information	128
 Chapter 12 Quality of Service (QoS)	
Overview of Policy-Based Quality of Service	133
Applications and Types of QoS	134
Voice Applications	134
Video Applications	134
Critical Database Applications	134
Web Browsing Applications	135
File Server Applications	135
Configuring QoS for a Port or VLAN	135
Traffic Groupings	136
Access List Based Traffic Groupings	136
MAC-Based Traffic Groupings	137
Explicit Class of Service (802.1p and DiffServ) Traffic Groupings	138
Configuring DiffServ	140
Physical and Logical Groupings	142
Verifying Configuration and Performance	143
QoS Monitor	143
Displaying QoS Profile Information	143
Modifying a QoS Configuration	144
Traffic Rate-Limiting	144
Dynamic Link Context System	144
DLCS Guidelines	145
DLCS Limitations	145
DLCS Commands	145
 Chapter 13 Status Monitoring and Statistics	
Status Monitoring	147
Port Statistics	149
Port Errors	149


Port Monitoring Display Keys	150
Setting the System Recovery Level	151
Logging	151
Local Logging	152
Remote Logging	153
Logging Configuration Changes	153
Logging Commands	153
RMON	155
About RMON	155
RMON Features of the Switch	155
Configuring RMON	156
Event Actions	157
 Chapter 14 Spanning Tree Protocol (STP)	
Overview of the Spanning Tree Protocol	159
Spanning Tree Domains	159
Defaults	160
STPD BPDU Tunneling	160
STP Configurations	160
Configuring STP on the Switch	162
STP Configuration Example	165
Displaying STP Settings	165
Disabling and Resetting STP	165
 Chapter 15 IP Unicast Routing	
Overview of IP Unicast Routing	167
Router Interfaces	168
Populating the Routing Table	168
Subnet-Directed Broadcast Forwarding	170
Proxy ARP	170
ARP-Incapable Devices	170
Proxy ARP Between Subnets	171
Relative Route Priorities	171
Configuring IP Unicast Routing	172
Verifying the IP Unicast Routing Configuration	172
IP Commands	172
Routing Configuration Example	177
Displaying Router Settings	178
Resetting and Disabling Router Settings	178

Configuring DHCP/BOOTP Relay	179
Verifying the DHCP/BOOTP Relay Configuration	180
UDP-Forwarding	180
Configuring UDP-Forwarding	180
UDP-Forwarding Example	181
ICMP Packet Processing	181
UDP-Forwarding Commands	181
 Chapter 16 Interior Gateway Routing Protocols	
Overview	183
RIP Versus OSPF	184
Overview of RIP	184
Routing Table	185
Split Horizon	185
Poison Reverse	185
Triggered Updates	185
Route Advertisement of VLANs	185
RIP Version 1 Versus RIP Version 2	185
Overview of OSPF	186
Link-State Database	186
Areas	187
Point-to-Point Support	190
Route Re-Distribution	191
Configuring Route Re-Distribution	191
OSPF Timers and Authentication	192
Configuring RIP	192
RIP Configuration Example	195
Displaying RIP Settings	196
Resetting and Disabling RIP	196
Configuring OSPF	197
Configuring OSPF Wait Interval	201
Displaying OSPF Settings	202
OSPF LSD Display	202
Resetting and Disabling OSPF Settings	203
 Chapter 17 IP Multicast Routing	
Overview	205
IGMP Overview	205
Configuring IP Multicasting Routing	206
Displaying IP Multicast Routing Settings	207

	Deleting and Resetting IP Multicast Settings	208
Appendix A	Safety Information	
	Important Safety Information	209
	Power	209
	Power Cord	210
	Connections	210
	Lithium Battery	210
Appendix B	Technical Specifications	
Appendix C	Supported Standards	
Appendix D	Software Upgrade and Boot Options	
	Downloading a New Image	217
	Rebooting the Switch	218
	Saving Configuration Changes	219
	Returning to Factory Defaults	219
	Using TFTP to Upload the Configuration	219
	Using TFTP to Download the Configuration	221
	Downloading a Complete Configuration	221
	Downloading an Incremental Configuration	221
	Scheduled Incremental Configuration Download	222
	Remember to Save	222
	Upgrading and Accessing BootROM	222
	Upgrading BootROM	222
	Accessing the BootROM menu	223
	Boot Option Commands	223
Appendix E	Troubleshooting	
	LEDs	225
	Using the Command-Line Interface	226
	Port Configuration	227
	VLANs	228
	STP	229
	Debug Tracing	229
	TOP Command	229
	Contacting Extreme Technical Support	229


Index

Index of Commands



Figures


1	Summit24e3 switch front view	20
2	Summit24e3 switch rear view	21
3	Fitting the mounting bracket	26
4	Null-modem cable pin-outs	28
5	PC-AT serial null-modem cable pin-outs	28
6	Mini-GBIC modules	31
7	Example of a port-based VLAN on the Summit24e3 switch	80
8	Single port-based VLAN spanning two switches	81
9	Two port-based VLANs spanning two switches	81
10	Physical diagram of tagged and untagged traffic	83
11	Logical diagram of tagged and untagged traffic	83
12	Permit-established access list example topology	100
13	Access control list denies all TCP and UDP traffic	101
14	Access list allows TCP traffic	102
15	Host A initiates a TCP session to host B	102
16	Permit-established access list filters out SYN packet to destination	103
17	ICMP packets are filtered out	103
18	RIP access policy example	106
19	OSPF access policy example	108
20	NAT Overview	111
21	Gigabit Ethernet fiber EAPS MAN ring	120
22	EAPS operation	121
23	EAPS fault detection and protection switching	122
24	Ethernet packet encapsulation	138
25	IP packet header encapsulation	140
26	Multiple Spanning Tree Domains	161
27	Tag-based STP configuration	162
28	Routing between VLANs	168
29	Unicast routing configuration example	177
30	Virtual link using Area 1 as a transit area	189
31	Virtual link providing redundancy	190
32	Route re-distribution	191
33	RIP configuration example	195



Tables

1	Notice Icons	17
2	Text Conventions	18
3	Summit24e3 Switch LED Behavior	21
4	Mini-GBIC Types and Distances	23
5	1000BASE-SX Specifications	23
6	Console Connector Pinouts	27
7	ExtremeWare Summit24e3 Factory Defaults	38
8	Command Syntax Symbols	41
9	Line-Editing Keys	41
10	Common Commands	42
11	Default Accounts	45
12	DNS Commands	47
13	Ping Command Parameters	48
14	SNMP Configuration Commands	55
15	RADIUS Commands	57
16	TACACS+ Commands	61
17	Network Login Configuration Commands	66
18	Greenwich Mean Time Offsets	68
19	SNTP Configuration Commands	69
20	Switch Port Commands	72
21	Switch Port-Mirroring Configuration Commands	76
22	EDP Commands	77
23	VLAN Configuration Commands	85
24	FDB Configuration Commands	89
25	Access Control List Configuration Commands	97
26	Routing Access Policy Configuration Commands	109
27	NAT Configuration Commands	114
28	NAT Timeout Commands	116
29	EAPS Commands	123
30	show eaps Display Fields	129
31	Traffic Type and QoS Guidelines	135
32	QoS Configuration Commands	135
33	Traffic Groupings by Precedence	136

34	802.1p Priority Value-to-QoS Profile to Hardware Queue Default Mapping	139
35	802.1p Configuration Commands	139
36	DiffServ Configuration Commands	140
37	Default Code Point-to-QoS Profile Mapping	141
38	DLCS Configuration Commands	145
39	Status Monitoring Commands	147
40	Port Monitoring Display Keys	150
41	Fault Levels Assigned by the Switch	151
42	Fault Log Subsystems	152
43	Logging Commands	153
44	Event Actions	157
45	STP Configuration Commands	163
46	STP Disable and Reset Commands	165
47	Relative Route Priorities	171
48	Basic IP Commands	173
49	Route Table Configuration Commands	174
50	ICMP Configuration Commands	175
51	Router Show Commands	178
52	Router Reset and Disable Commands	178
53	UDP-Forwarding Commands	181
54	LSA Type Numbers	186
55	OSPF Link Types	190
56	RIP Configuration Commands	192
57	RIP Show Commands	196
58	RIP Reset and Disable Commands	196
59	OSPF Configuration Commands	197
60	OSPF Show Commands	202
61	OSPF Reset and Disable Commands	203
62	IGMP Configuration Commands	207
63	IP Multicast Routing Show Commands	207
64	IP Multicast Routing Reset and Disable Commands	208
65	Boot Option Commands	223



Preface

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

Introduction

This guide provides the required information to install the Summit24e3 switch and configure the ExtremeWare™ software running on the Summit24e3 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
- 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.

Table 1: Notice Icons (continued)

Icon	Notice Type	Alerts you to...
	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 <i>italicized</i> 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 Release Notes*
- *Summit24e3 Switch Release Notes*

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

- <http://www.extremenetworks.com/>



Summit24e3 Switch Overview

This chapter describes the features and functionality of the Summit24e3 switch:

- Summary of Features on page 19
- Summit24e3 Switch Front View on page 20
- Summit24e3 Switch Rear View on page 21
- Summit24e3 Switch LEDs on page 21
- Mini-GBIC Type and Hardware/ Software Support on page 22

Summary of Features

The Summit24e3 supports the following ExtremeWare features:


- Virtual local area networks (VLANs) including support for IEEE 802.1Q and IEEE 802.1p
- Spanning Tree Protocol (STP) (IEEE 802.1D)
- Quality of Service (QoS) including support for IEEE 802.1p, MAC QoS, and four hardware queues
- Wire-speed Internet Protocol (IP) routing
- DHCP/BOOTP Relay
- Network Address Translation (NAT)
- Extreme Standby Router Protocol (ESRP) - Aware support
- Ethernet Automated Protection Switching (EAPS) support
- 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
- Load sharing on multiple ports

- RADIUS client and per-command authentication support
- TACACS+ support
- Network Login
- Console command-line interface (CLI) connection
- Telnet CLI connection
- SSH2 connection
- Simple Network Management Protocol (SNMP) support
- Remote Monitoring (RMON)
- Traffic mirroring for ports

Summit24e3 Switch Front View

Figure 1 shows the Summit24e3 switch front view.

Figure 1: Summit24e3 switch front view



The Summit24e3 switch has 24 autosensing 10BASE-T/100BASE-TX ports using RJ-45 connectors and provides two unpopulated Extreme mini-GBIC ports using LC connectors. You can use the 1000BASE-SX mini-GBIC in the Summit24e3 switch.

Console Port

Use the console port (9-pin, “D” type connector) for connecting a terminal and carrying out local management.

Reset Button

The reset button is used to hard reset the switch. Use a non-conductive tool to push the reset button.



NOTE

See Table 4 for information about supported mini-GBIC types and distances.


**NOTE**

See “Summit24e3 Switch LEDs” on page 21 for more details.

Summit24e3 Switch Rear View

Figure 2 shows the rear view of the Summit24e3 switch.

Figure 2: Summit24e3 switch rear view



Power Socket

The Summit24e3 switch automatically adjusts to the supply voltage. The power supply operates down to 90 V.

Serial Number

Use this serial number for fault-reporting purposes.

MAC Address

This label shows the unique Ethernet MAC address assigned to this device.

**NOTE**

The Summit24e3 switch certification and safety label is located on the bottom of the switch.

Summit24e3 Switch LEDs

Table 3 describes the light emitting diode (LED) behavior on the Summit24e3 switch.

Table 3: Summit24e3 Switch LED Behavior

Unit Status LED (MGMT LED)

Color	Indicates
Green blinking	The Summit switch is operating normally.
Amber	The Summit has failed its POST.

Table 3: Summit24e3 Switch LED Behavior (continued)**10/100 Mbps Port Status LEDs**

Color	Indicates
Green	Link is present; port is enabled.
Green blinking	Link is present; there is activity.
Off	Link is not present or the port is disabled.

Mini-GBIC Port Status LEDs

Color	Indicates
Amber	Frames are being transmitted/received on this port.
Amber blinking	A non-supported GBIC is installed on this port.
Green	Link is present; port is enabled; full-duplex operation.
Green blinking	Link is present; port is disabled.
Off	Link is not present.

Port Connections

The Summit24e3 switch uses a combination of 10BASE-T/100BASE-TX ports using RJ-45 connectors and Small Form Factor Pluggable (SFP) Gigabit Interface Connectors (GBICs), also known as mini-GBICs, using LC connectors and supports 1000BASE-SX.

Full-Duplex

The Summit24e3 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 Mbps ports on the Summit24e3 switch autonegotiates for half- or full-duplex operation.

Mini-GBIC Type and Hardware/Software Support

The Summit24e3 switch supports the SFP GBIC, also known as the mini-GBIC. The system uses identifier bits to determine the media type of the mini-GBIC that is installed. The Summit24e3 supports only the mini-GBIC.

This section describes the mini-GBIC type and specifications.

Mini-GBIC Type and Specifications

Table 4 describes the mini-GBIC type and distances for the Summit24e3 switch.

Table 4: Mini-GBIC Types and Distances

Standard	Media Type	Mhz•Km Rating	Maximum Distance (Meters)
1000BASE-SX (850 nm optical window)	50/125 μ m multimode fiber	400	500
	50/125 μ m multimode fiber	500	550
	62.5/125 μ m multimode fiber	160	220
	62.5/125 μ m multimode fiber	200	275

Table 5 describes the specifications for the 1000BASE-SX Mini-GBIC.

Table 5: 1000BASE-SX 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

Safety Information

Before you begin the process of installing or replacing a mini-GBIC, read the safety information in this section.



CAUTION

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

Mini-GBICs are a class 1 laser device. Use only Extreme-approved devices.



NOTE

If you see an amber blinking mini-GBIC port status LED on your Summit24e3 switch, you do not have an Extreme-supported mini-GBIC installed in your switch. To correct this problem, make sure you install an Extreme-supported mini-GBIC into the port on the switch.

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

2

Switch Installation

This chapter describes the following topics:

- Determining the Switch Location on page 25
- Installing the Switch on page 26
- Connecting Equipment to the Console Port on page 27
- Powering On the Switch on page 28
- Checking the Installation on page 28
- Logging In for the First Time on page 29
- Installing or Replacing a Mini-Gigabit Interface Connector (Mini-GBIC) on page 30



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 Summit24e3 switch 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 25 mm (1 inch) 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 Appendix A of this guide.

Installing the Switch

The Summit24e3 switch 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 Summit24e3 switch:

- 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 3.

Figure 3: Fitting the mounting bracket



SH_24e3rack

- 5 Repeat steps 2 - 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).
- 9 Connect cables.

Free-Standing

The Summit24e3 switch 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.

Stacking the Switch and Other Devices

You can place up to four Summit switches on top of one another.



NOTE

This relates only to physically placing the devices 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.

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 6.

Table 6: 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

Table 6: Console Connector Pinouts

Function	Pin Number	Direction
RTS (request to send)	7	Out
CTS (clear to send	8	In

Figure 4 shows the pin-outs for a 9-pin to RS-232 25-pin null-modem cable.

Figure 4: Null-modem cable pin-outs





Figure 5 shows the pin-outs for a 9-pin to 9-pin PC-AT null-modem serial cable.

Figure 5: 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. Turn the on/off switch to the on position.

Checking the Installation

After turning on power to the Summit24e3 switch, 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. The MGMT LED flashes until the switch successfully passes the POST.

If the switch passes the POST, the MGMT LED blinks at a slow rate (1 blink per second). If the switch fails the POST, the MGMT LED shows a solid amber light.

**NOTE**

For more information on the LEDs, see Chapter 1, “Summit24e3 Switch Overview”.

Logging In for the First Time

After the Summit24e3 switch 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 manually configure the IP settings, perform the following 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.

**NOTE**

For more information on switch security, see Chapter 4, “Accessing the Switch”.

- 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, *Summit24e3*) 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 typing
save

**NOTE**

For more information on saving configuration changes, refer to the ExtremeWare Software User Guide.

- 7 When you are finished using the facility, logout of the switch by typing
logout

**NOTE**

After two incorrect login attempts, the Summit24e3 switch locks you out of the login facility. You must wait a few minutes before attempting to log in again.

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 begin the process of installing or replacing a mini-GBIC, read the safety information in this section.

**CAUTION**

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

Mini-GBICs are a class 1 laser device. Use only Extreme-approved devices.

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.

Installing and Removing a Mini-GBIC

You can add and remove mini-GBICs from your Summit24e3 switch without powering off the system. Figure 6 shows the two types of mini-GBIC connectors.

Figure 6: Mini-GBIC modules



Mini-GBICs are a Class 1 laser device. Use only Extreme-approved devices.



NOTE

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



CAUTION

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



NOTE

If you see an amber blinking Mini-GBIC port status LED on your Summit24e3 switch, you do not have an Extreme-supported mini-GBIC installed in your switch. To correct this problem, make sure you install an Extreme-supported mini-GBIC into the port on the switch.

To remove the mini-GBIC connector labeled “Module A,” gently squeeze the sides to release it, and pull the mini-GBIC out of the slot.

To remove the mini-GBIC connector labeled “Module B,” lift down on the front handle and pull the mini-GBIC out of the slot.

To insert a mini-GBIC connector:

- 1** Holding the mini-GBIC by its sides, insert the mini-GBIC into the port on the switch. Mini-GBICs can be installed in ports 25 and 26 on the Summit24e3 switch.
- 2** Slide the mini-GBIC into the port, until you hear it click. If the mini-GBIC has a handle, push up on the handle to secure the mini-GBIC.



ExtremeWare Overview

This chapter describes the following topics:

- Summary of Features on page 33
- Security Licensing on page 37
- Software Factory Defaults on page 38

ExtremeWare is the full-featured software operating system that is designed to run on the Summit24e3 switch. This section describes the supported ExtremeWare features for the Summit24e3 switch.

Summary of Features

The Summit24e3 supports the following ExtremeWare features:

- Virtual local area networks (VLANs) including support for IEEE 802.1Q and IEEE 802.1p
- Spanning Tree Protocol (STP) (IEEE 802.1D)
- Quality of Service (QoS) including support for IEEE 802.1p, MAC QoS, and four hardware queues
- Wire-speed Internet Protocol (IP) routing
- DHCP/BOOTP Relay
- Network Address Translation (NAT)
- Extreme Standby Router Protocol (ESRP) - Aware support
- Ethernet Automated Protection Switching (EAPS) support
- 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
- Load sharing on multiple ports

- RADIUS client and per-command authentication support
- TACACS+ support
- Network Login
- Console command-line interface (CLI) connection
- Telnet CLI connection
- SSH2 connection
- Simple Network Management Protocol (SNMP) support
- Remote Monitoring (RMON)
- Traffic mirroring for ports

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:

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



For more information on VLANs, see Chapter 7, “Virtual LANs (VLANs)”.

Spanning Tree Protocol

The Summit24e3 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 14, “Spanning Tree Protocol (STP)”.

Quality of Service

ExtremeWare has Quality of Service (QoS) features that support IEEE 802.1p, MAC QoS, and four queues. These features 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 rate-limiting access control lists and apply them to different traffic types so that they have different maximum bandwidth, and priority.



NOTE

For more information on Quality of Service, see Chapter 12, “Quality of Service (QoS)”.

Unicast Routing

The Summit24e3 can route IP traffic between the VLANs that are configured as virtual router interfaces. Static IP routes are maintained in the routing table. The following routing protocols are supported:

- RIP version 1
- RIP version 2
- OSPF



NOTE

For more information on IP unicast routing, see Chapter 15, “IP Unicast Routing”.

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 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.



NOTE

For information on load sharing, see Chapter 6, “Configuring Ports on a Switch”.

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.

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 above), 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 using the protocol filter `any`. ESRP will not function correctly if the ESRP-aware switch interconnection port is configured for a protocol-sensitive VLAN using untagged traffic.

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.

Feature Licensing

Summit24e3 switches support software licensing for different levels of functionality. In ExtremeWare version 6.2e.1 (release 2.0 for the Summit24e3 platform), and above, feature 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*. Summit24e3 switches have Edge functionality, without the requirement of a license key. Edge functionality includes all switching functions, and also includes all available layer 3 QoS, access list, and ESRP-aware functions. Layer 3 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

Advanced Edge Functionality

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

- Rate-limiting ACLs
- IP routing using OSPF
- EAPS Edge (cannot be a core node on the ring)
- Network Login
- RADIUS
- TACACS+
- NAT

Verifying the Router License

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

Obtaining a Router License

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 switch arrives packaged with a certificate that contains the unique license key(s), and instructions for enabling the correct functionality on the switch. The certificate is typically packaged with the switch documentation. Once the license key is entered, it should not be necessary to enter the information again. However, we recommend keeping the certificate for your records.

You can upgrade the router 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

ExtremeWare version 6.0 and above supports the SSH2 protocol. SSH2 allows the encryption of Telnet session data. The encryption methods used are under U.S. export restriction control.

Software Factory Defaults

Table 7 shows factory defaults for Summit24e3 ExtremeWare features.

Table 7: ExtremeWare Summit24e3 Factory Defaults

Item	Default Setting
Serial or Telnet user account	<i>admin</i> with no password and <i>user</i> with no password
Telnet	Enabled
SSH2	Disabled
SNMP	Enabled
SNMP read community string	<i>public</i>
SNMP write community string	<i>private</i>
RMON	Disabled
BOOTP	Enabled on the default VLAN (<i>default</i>)
QoS	All traffic is part of the default queue
802.1p priority	Recognition enabled
802.3x flow control	Enabled on Gigabit Ethernet ports
Virtual LANs	Two VLANs predefined. VLAN named <i>default</i> contains all ports and belongs to the STPD named <i>s0</i> .
802.1Q tagging	All packets are untagged on the default VLAN (<i>default</i>).
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
NTP	Disabled
DNS	Disabled
EAPS	Disabled
NAT	Disabled
Network Login	Disabled
RADIUS	Disabled
TACACS+	Disabled
Port Mirroring	Disabled



NOTE

For default settings of individual ExtremeWare features, see the applicable individual chapters in this guide.

4

Accessing the Switch

This chapter describes the following topics:

- Understanding the Command Syntax on page 39
- Line-Editing Keys on page 41
- Command History on page 42
- Common Commands on page 42
- Configuring Management Access on page 44
- Domain Name Service Client Services on page 47
- Checking Basic Connectivity on page 47

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.

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 D, “Software Upgrade and Boot Options”.*

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 [Return]. The syntax helper provides a list of options for the remainder of the command.

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

Command Completion with Syntax Helper

ExtremeWare provides command completion by way of the [Tab] key. If you enter a partial command, pressing the [Tab] key posts a list of available options, and places the cursor at the end of the command.

Abbreviated Syntax

Abbreviated syntax is the most unambiguous, shortest allowable abbreviation of a command or parameter. Typically, this is the first three letters of the command.

In command tables throughout this guide, abbreviated syntax is noted using bold characters.



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 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, on the stand-alone switch, instead of entering the command

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

you could enter the following shortcut:

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

Summit24e3 Switch Numerical Ranges

Commands that require you to enter one or more port numbers on a Summit24e3 switch 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.

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 8 summarizes command syntax symbols.

Table 8: Command Syntax Symbols

angle brackets < >	Enclose a variable or value. You must specify the variable or value. For example, in the syntax <pre>config vlan <name> ipaddress <ip_address></pre> you must supply a VLAN name for <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 <pre>use image [primary secondary]</pre> 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 <pre>config snmp community [read-only read-write] <string></pre> 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 <pre>reboot {<date> <time> cancel}</pre> you can specify either a particular date and time combination, or the keyword <code>cancel</code> 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.

Line-Editing Keys

Table 9 describes the line-editing keys available using the CLI.

Table 9: Line-Editing Keys

Symbol	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.

Table 9: Line-Editing Keys (continued)

Symbol	Description
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 moves 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.

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 10 describes common commands used to manage the switch. Commands specific to a particular feature are described in the other chapters of this guide.

Table 10: Common Commands

Command	Description
clear session <number>	Terminates a Telnet session from the switch.
config account <username> {encrypted} {<password>}	Configures a user account password. Passwords must have a minimum of 1 character and can have a maximum of 32 characters. User names and passwords are case-sensitive.
config 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.
config ports <portlist> 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.
config ssh2 key {pregenerated}	Generates the SSH2 host key.

Table 10: Common Commands (continued)

Command	Description
config sys-recovery-level [none critical all]	Configures a recovery option for instances where an exception occurs in ExtremeWare. Specify one of the following: <code>none</code> — Recovery without system reboot. <code>critical</code> — ExtremeWare logs an error to the syslog, and reboots the system after critical exceptions. <code>all</code> — ExtremeWare logs an error to the syslog, and reboots the system after any exception. The default setting is <code>none</code>.
config time <date> <time>	Configures the system date and time. The format is as follows: <pre>mm/dd/yyyy hh:mm:ss</pre> The time uses a 24-hour clock format. You cannot set the year past 2036.
config timezone <gmt_offset> {autodst noautodst}	Configures the time zone information to the configured offset from GMT time. The format of <code>gmt_offset</code> is +/- minutes from GMT time. Specify: <code>autodst</code> — Enables automatic Daylight Savings Time change. <code>noautodst</code> — Disables automatic Daylight Savings Time change. The default setting is <code>autodst</code>.
config vlan <name> ipaddress <ip_address> {<mask>}	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 32 characters, the password is between 0 and 16 characters.
create vlan <name>	Creates a VLAN.
delete account <username>	Deletes a user account.
delete vlan <name>	Deletes a VLAN.
disable bootp 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>	Disables a port on the switch.

Table 10: Common Commands (continued)

Command	Description
disable ssh2	Disables SSH2 Telnet access to the switch.
disable telnet	Disables Telnet access to the switch.
enable bootp 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 <code>show</code> 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 ssh2 {port <tcp_port_number>}	Enables SSH2 Telnet sessions. By default, SSH2 uses TCP port number 22.
enable telnet {port <tcp_port_number>}	Enables Telnet access to the switch. By default, Telnet uses TCP port number 23.
history	Displays the previous 49 commands entered on the switch.
show banner	Displays the user-configured banner.
unconfig 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 <code>all</code> , 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 5, “Managing the Switch”.

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:

```
Summit24e3:2>
```

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:

```
Summit24e3:18#
```

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:

```
*Summit24e3:19#
```

Default Accounts

By default, the switch is configured with two accounts, as shown in Table 11.

Table 11: 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 must have a minimum of four characters and can have a maximum of 12 characters.



NOTE

User names and passwords are 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:

```
config 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:

```
config 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 31 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 | 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

The account name admin cannot be deleted.

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.

Table 12 describes the commands used to configure DNS.

Table 12: DNS Commands

Command	Description
config dns-client add <ipaddress>	Adds a DNS name server(s) to the available server list for the DNS client. Up to three name servers can be configured.
config dns-client default-domain <domain_name>	Configures the domain that the DNS client uses if a fully qualified domain name is not entered. For example, if the default domain is configured to be foo.com, executing ping bar searches for bar.foo.com.
config dns-client delete <ipaddress>	Removes a DNS server.
nslookup <hostname>	Displays the IP address of the requested host.
show dns-client	Displays the DNS configuration.

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

Options for the `ping` command are described in Table 13.

Table 13: Ping Command Parameters

Parameter	Description
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 <code>start_size</code> and <code>end_size</code> are specified, transmits ICMP requests using 1 byte increments, per packet. If no <code>end_size</code> is specified, packets of <code>start_size</code> are sent.
<ipaddress>	Specifies the IP address of the host.
<hostname>	Specifies the name of the host. To use the <code>hostname</code> , 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.

Traceroute

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

```
traceroute [<ip_address> | <hostname>] {from <src_ipaddress>} {ttl <TTL>} {port
<port>}
```

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 up to the time-to-live number of the switch.
- `port` uses the specified UDP port number.

5

Managing the Switch

This chapter describes the following topics:

- Overview on page 49
- Using the Console Interface on page 50
- Using Telnet on page 50
- Using Secure Shell 2 (SSH2) on page 53
- Using SNMP on page 54
- Authenticating Users on page 56
- Using Network Login on page 62
- Using the Simple Network Time Protocol on page 67

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. Remote access includes:
 - Telnet using the CLI interface.
 - SSH2 using the CLI interface.
 - SNMP access using ExtremeWare Enterprise Manager or another SNMP manager.

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

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

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 front of the Summit 24e3 switch.

Once the connection is established, you will see the switch prompt and you can log in.

Using Telnet

Any workstation with a Telnet facility should be able to communicate with the switch over a TCP/IP network.

Up to eight active Telnet sessions can access the switch concurrently. If `idle timeouts` 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 configure the switch IP parameters. See “Configuring Switch IP Parameters” on page 50 for more information. Telnet is enabled by default.

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.

Once 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 add 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)

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

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

```
enable bootp 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 must be capable of differentiating its relay based on the gateway portion of the BOOTP packet.



NOTE

For more information on DHCP/BOOTP relay, see Chapter 15, “IP Unicast Routing”.

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.
- 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 a VLAN. The switch can be assigned multiple IP addresses.



NOTE

For information on creating and configuring VLANs, see Chapter 7, “Virtual LANs (VLANs)”.

To manually configure the IP settings, 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 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:

```
config vlan <name> ipaddress <ipaddress> {<subnet_mask>}
```

For example:

```
config 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:

```
config vlan default ipaddress 123.45.67.8 / 24
```

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

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

For example:

```
config 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 typing:

```
save
```

- 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 <session_number>
```

Controlling Telnet Access

By default, Telnet services are enabled on the switch. 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.

Using Secure Shell 2 (SSH2)

Secure Shell 2 (SSH2) is a feature of ExtremeWare that allows you to encrypt Telnet session data between the switch and a network administrator using SSH2 client software. 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, refer to 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.

Enabling SSH2

Because SSH2 is currently under U.S. export restrictions, before enabling SSH2, you must first obtain a security license from Extreme Networks. The procedure for obtaining a security license key is described in Chapter 3, “ExtremeWare Overview”.

To enable SSH2, use the following command:

```
enable ssh2 {port <tcp_port_number>}
```

An authentication key must be generated for each SSH2 session. This can be done automatically by the switch or by the client application. To have the key generated by the switch, use the following command:

```
config ssh2 key {pregenerated}
```

If you do not select automatic key generation, you are prompted to enter the key when you enable SSH2.

You can specify a TCP port number to be used for SSH2 communication. By default the TCP port number is 22.

The supported cipher is 3DES-CBC. The supported key exchange is DSA.

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>.

After you obtain the SSH2 key value, copy the key to the SSH2 client application. Also, ensure that the client is configured for any nondefault TCP port information that you have configured on the switch. Once these tasks are accomplished, you may form an SSH2-encrypted session with the switch.

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.

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.

Supported MIBs

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

Configuring SNMP Settings

The following SNMP 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 SNMP traps to all trap receivers. You can have a maximum of 16 trap receivers configured for each switch. Entries in this list can also be created, modified, and deleted using the RMON2 trapDestTable MIB variable, as described in RFC 2021.
- **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*. A total of eight community strings can be configured on the switch. The community string for all authorized trap receivers must be configured on the switch for the trap receiver to receive switch-generated traps. SNMP community strings can contain up to 127 characters.
- **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.

Table 14 describes SNMP configuration commands.

Table 14: SNMP Configuration Commands

Command	Description
config snmp add trapreceiver <ipaddress> community <string>	Adds the IP address of a specified trap receiver. The IP address can be a unicast, multicast, or broadcast address. A maximum of 16 trap receivers is allowed.
config snmp community [read-only read-write] <string>	Adds an SNMP read or read/write community string. The default read-only community string is <code>public</code> . The default read-write community string is <code>private</code> . Each community string can have a maximum of 127 characters, and can be enclosed by double quotation marks.
config snmp delete trapreceiver [<ip_address> community <string> all]	Deletes the IP address of a specified trap receiver or all authorized trap receivers.
config snmp syscontact <string>	Configures the name of the system contact. A maximum of 255 characters is allowed.
config snmp syslocation <string>	Configures the location of the switch. A maximum of 255 characters is allowed.
config snmp sysname <string>	Configures the name of the switch. A maximum of 32 characters is allowed. The default sysname is the model name of the device (for example, <code>Summit24e3</code>). The <code>sysname</code> appears in the switch prompt.
disable snmp access	Disables SNMP on the switch. Disabling SNMP access does not affect the SNMP configuration (for example, community strings).
disable snmp traps	Prevents SNMP traps from being sent from the switch. Does not clear the SNMP trap receivers that have been configured.
enable snmp access	Turns on SNMP support for the switch.
enable snmp traps	Turns on SNMP trap support.
unconfig management	Restores default values to all SNMP-related entries.

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, and SNMP
- SNMP community strings
- Authorized SNMP station list
- SNMP trap receiver list
- RMON polling configuration
- Login statistics

Authenticating Users

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 or console access to the switch.



You cannot configure RADUIS 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 an 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.

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 “Configuring RADIUS Client” on page 56.

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 commands are described in Table 15.

Table 15: RADIUS Commands

Command	Description
<pre>config radius [primary secondary] server [<ipaddress> <hostname>] {<udp_port>} client-ip <ipaddress></pre>	Configures the primary and secondary RADIUS server. Specify the following: • [primary secondary] — Configure either the primary or secondary RADIUS server. • [<ipaddress> <hostname>] — The IP address or hostname of the server being configured. • <udp_port> — The UDP port to use to contact the RADIUS server. The default UDP port setting is 1645. • client-ip <ipaddress> — The IP address used by the switch to identify itself when communicating with the RADIUS server. The RADIUS server defined by this command is used for user name authentication and CLI command authentication.
<pre>config radius [primary secondary] shared-secret {encrypted} <string></pre>	Configures the authentication string used to communicate with the RADIUS server.
<pre>config radius-accounting [primary secondary] server [<ipaddress> <hostname>] {<udp_port>} client-ip <ipaddress></pre>	Configures the RADIUS accounting server. Specify the following: • [primary secondary] — Configure either the primary or secondary RADIUS server. • [<ipaddress> <hostname>] — The IP address or hostname of the server being configured. • <udp_port> — The UDP port to use to contact the RADIUS server. The default UDP port setting is 1646. • client-ip <ipaddress> — The IP address used by the switch to identify itself when communicating with the RADIUS server. The accounting server and the RADIUS authentication server can be the same.
<pre>config radius-accounting [primary secondary] shared-secret {encrypted} <string></pre>	Configures the authentication string used to communicate with the RADIUS accounting server.
<pre>disable radius</pre>	Disables the RADIUS client.
<pre>disable radius-accounting</pre>	Disables RADIUS accounting.
<pre>enable radius</pre>	Enables the RADIUS client. When enabled, all CLI logins are sent to the RADIUS servers for authentication. When used with a RADIUS server that supports ExtremeWare CLI authorization, each CLI command is sent to the RADIUS server for authentication before it is executed.

Table 15: RADIUS Commands (continued)

Command	Description
enable radius-accounting	Enables RADIUS accounting. The RADIUS client must also be enabled.
show radius	Displays the current RADIUS client configuration and statistics.
show radius-accounting	Displays the current RADIUS accounting client configuration and statistics
unconfig radius {server [primary secondary]}	Unconfigures the radius client configuration.
unconfig radius-accounting {server [primary secondary]}	Unconfigures the radius accounting client configuration.

RADIUS RFC 2138 Attributes

The RADIUS RFC 2138 optional attributes supported are as follows:

- User-Name
- User-Password
- Service-Type
- Login-IP-Host

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 web server at <http://www.extremenetworks.com/extreme/support/otherapps.htm> or 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 counter 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.

Table 16 describes the commands that are used to configure TACACS+.

Table 16: TACACS+ Commands

Command	Description
config tacacs [primary secondary] server [<ipaddress> <hostname>] {<udp_port>} client-ip <ipaddress>	Configure the server information for a TACACS+ server. Specify the following: - primary secondary — Specifies primary or secondary server configuration. To remove a server, use the address 0.0.0.0. <ipaddress> <hostname> — Specifies the TACACS+ server. <udp_port> — Optionally specifies the UDP port to be used. - client-ip — Specifies the IP address used by the switch to identify itself when communicating with the TACACS+ server.
config tacacs [primary secondary] shared-secret {encrypted} <string>	Configures the shared secret string used to communicate with the TACACS+ server.
config tacacs-accounting [primary secondary] server [<ipaddress> <hostname>] {<udp_port>} client-ip <ipaddress>	Configures the TACACS+ accounting server. You can use the same server for accounting and authentication.
config tacacs-accounting [primary secondary] shared-secret {encrypted} <string>	Configures the shared secret string used to communicate with the TACACS+ accounting server.
disable tacacs	Disables TACACS+.
disable tacacs-accounting	Disables TACACS+ accounting.
disable tacacs-authorization	Disables CLI command authorization.
enable tacacs	Enables TACACS+. Once enabled, all CLI logins are sent to one of the two TACACS+ server for login name authentication and accounting.
enable tacacs-accounting	Enables TACACS+ accounting. If accounting is use, the TACACS+ client must also be enabled.
enable tacacs-authorization	Enables CLI command authorization. When enabled, each command is transmitted to the remote TACACS+ server for authorization before the command is executed.

Table 16: TACACS+ Commands (continued)

Command	Description
show tacacs	Displays the current TACACS+ configuration and statistics.
show tacacs-accounting	Displays the current TACACS+ accounting client configuration and statistics.
unconfig tacacs {server [primary secondary]}	Unconfigures the TACACS+ client configuration.
unconfig tacacs-accounting {server [primary secondary]}	Unconfigures the TACACS+ accounting client configuration.

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.



NOTE

Windows authentication is not supported via network login.

Network login has two modes of operation:

- Campus mode

Campus mode is used when a port in a VLAN will move to another VLAN when authentication has been completed successfully. This mode is for the roaming user who will not always be using the same port for authentication.

- ISP mode

ISP mode is used when the port and VLAN used will remain constant. All network settings are configured for that VLAN.

These two network login modes have the following functional similarities:

- Until authentication takes place, ports on the VLAN are kept in a non-forwarding state.
- each mode requires the user to open a web browser with the IP address of the switch. This is the only address that the client can reach in a non-authenticated state.
- The web server on the switch provides user authentication.
- After authentication takes place, ports are moved into a forwarding state and moved to the VLAN configuration on the RADIUS server.

Using Network Login in Campus Mode

Campus mode requires:

- A DHCP server
- A RADIUS server configuration

The RADIUS server must have the following options configured in its dictionary file for network login:

Extreme.attr Extreme-Netlogin-Vlan 203 string (1, 0, ENCAPS)

The following optional configuration parameters can also be specified:

Extreme.attr Extreme-Netlogin-Url 204 string (1, 0, ENCAPS)

Extreme.attr Extreme-Netlogin-Url-Desc 205 string (1, 0, ENCAPS)



NOTE

These settings are for the Merit 3.6 version of RADIUS. The syntax of these settings will vary based on the type of RADIUS server that you are using.

The RADIUS server must also contain entries in the user file for a permanent VLAN, the URL to be redirected to after authentication has taken place, and the description of that URL. For example:

```
auto    Authentication-Type = Unix-PW, Service-Type = login
        Filter-Id = "unlim"
        Extreme:Extreme-Netlogin-Vlan = "corp"
        Extreme:Extreme-Netlogin-Url = "http://192.207.37.16"
        Extreme:Extreme-Netlogin-Url-Desc = "Extreme Networks Home"
```

In this example, the username is *auto*, the permanent VLAN is *corp*, and the URL to be redirected to is the Extreme Networks home page *http://192.207.37.16*.

Configuring Campus Mode

To configure the switch to use network login in campus mode, follow these steps:

- 1 Configure the switch as a RADIUS client. See “RADIUS Client” on page 56.
- 2 Configure a DHCP range for the port or ports in the VLAN on which you want to enable network login, using this command:

```
config vlan <name> dhcp-address-range <ipaddress1> - <ipaddress2>
```

The switch will assign a temporary DHCP address within the DHCP range to the client.

- 3 Enable network login on the port, using the command:

```
enable netlogin ports <portlist> vlan <name>
```



NOTE

Network login is used on a per port, per VLAN basis. A port that is tagged can belong to more than one VLAN. In this case, network login can be enabled on one port for each VLAN.

Example Configuration Using Campus Mode

This example creates a permanent VLAN named *corp* on the switch. This VLAN will be used for authentication through a RADIUS server. The RADIUS server is 10.201.26.243 and the IP address of the switch is 10.201.26.11. The secret is “secret”. A temporary VLAN named *temporary* is created and port 9 is added. Network login is enabled on the port.

```
create vlan corp
```

```

config corp ipaddress 10.201.26.11/24
config radius primary server 10.201.26.243 client-ip 10.201.26.11
config radius primary shared-secret secret
enable radius
create vlan temporary
config temporary add port 9
config temporary ipaddress 192.168.0.1/24
config temporary dhcp-address-range 192.168.0.20 - 192.168.0.100
enable netlogin ports 9 vlan temporary

```

User Login Using Campus Mode

To log in as a user from the client, 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.
In this example, the user will plug into port 9.
- 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 IP address 192.168.0.20.

- 5 Bring up the web browser and enter the IP address of the switch.



NOTE

It is important to use the IP address of a VLAN that is reachable from anywhere on the network

A page will open 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 86.

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.
- 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 completed. The login process is completed when you receive a permanent address.

Using Network Login in ISP Mode

In ISP mode, a RADIUS server might be used to provide user authentication. No Extreme-specific lines are required for the dictionary or the user file.

Configuring ISP Mode

Configure the switch to use network login in ISP mode, using this command:

```
enable netlogin ports <portlist> vlan <name>
```



NOTE

Network login is used on a per port, per VLAN basis. A port that is tagged can belong to more than one VLAN. In this case, network login can be enabled on one port for each VLAN.

Example Configuration Using ISP Mode

This example creates a permanent VLAN named *corp* on the switch. This VLAN will be used for authentication through RADIUS. The radius server is 10.201.26.243 and the IP address of the switch is 10.201.26.11. The secret is “secret”. Port 9 is added to the VLAN *corp*. Network login is enabled on the port.

```
create vlan corp
config corp ipaddress 10.201.26.11/24
config radius primary server 10.201.26.243 client-ip 10.201.26.11
config radius primary shared-secret secret
enable radius
config corp add port 9
enable netlogin ports 9 vlan corp
```

DHCP Server on the Switch

A DHCP server with limited configuration capabilities is included in the switch to provide IP addresses to clients.

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 <name>
disable dhcp ports <portlist> vlan <name>
```

Network Login Configuration Commands

Table 17 describes the commands used to configure network login.

Table 17: Network Login Configuration Commands

Command	Description
config vlan <name> dhcp-address-range <ipaddress1> - <ipaddress2>	Configures a set of DHCP addresses for a VLAN.
config vlan <name> dhcp-lease-timer <lease-timer>	Configures the timer value in seconds returned as part of the DHCP response.
config vlan <name> dhcp-options [default-gateway dns-server wins-server] <ipaddress>	Configures the DHCP options returned as part of the DHCP response by a switch configured as a DHCP server.
config vlan <name> netlogin-lease-timer <lease-timer>	Configures the timer value in seconds returned as part of the DHCP response for clients attached to network enabled ports. The default value is 30 seconds.
disable dhcp ports <portlist> vlan <name>	Disables DHCP on a specified port in a VLAN.
disable netlogin ports <portlist> vlan <name>	Disables network login on a specified port in a VLAN.
enable dhcp ports <portlist> vlan <name>	Enables DHCP on a specified port in a VLAN.
enable netlogin ports <portlist> vlan <name>	Enables network login on a specified port in a VLAN.

Displaying Network Login Settings

To display the network login settings, use the following command:

```
show netlogin info {ports <portlist> vlan <name>}
```

Example

```
#show netlogin info ports 9 vlan temporary
Port 9: VLAN: temporary
Port State: Not Authenticated
Temp IP: Unknown
DHCP: Not Enabled
User: Unknown MAC: Unknown
```

In this example, the user is using campus mode and no authentication has taken place. Therefore, the port state displays as not authenticated. No packets sent by the user on port 9 will get past the port

until authentication takes place. After authentication has taken place and the permanent IP address is obtained, the show command displays the port state as authenticated.

```
#show netlogin info ports 9 vlan corp
Port 9: VLAN: corp
Port State: Authenticated
Temp IP: Unknown
DHCP: Not Enabled
User: auto MAC: 00:10:A4:A9:11:3B
```

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 <name>
```

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 Savings 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 Savings Time preference. The command syntax to configure GMT offset and usage of Daylight Savings is as follows:

```
config timezone <GMT_offset> {autodst | noautodst}
```

The GMT_OFFSET is in +/- minutes from the GMT time. Automatic Daylight Savings Time (DST) changes can be enabled or disabled. The default setting is enabled.

- 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:

```
config sntp-client [primary | secondary] server [<ip_address> | <hostname>]
```

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:

```
config 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, Daylight Savings Time, 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 18 describes GMT offsets.

Table 18: Greenwich Mean Time Offsets

GMT Offset in Hours	GMT Offset in Minutes	Common Time Zone References	Cities
+0:00	+0	GMT - Greenwich Mean UT or UTC - Universal (Coordinated) WET - Western European	London, England; Dublin, Ireland; Edinburgh, Scotland; Lisbon, Portugal; Reykjavik, Iceland; Casablanca, Morocco
-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	

Table 18: Greenwich Mean Time Offsets (continued)

GMT Offset in Hours	GMT Offset in Minutes	Common Time Zone References	Cities
+1:00	+60	CET - Central European FWT - French Winter MET - Middle European MEWT - Middle European Winter SWT - Swedish Winter	Paris, France; Berlin, Germany; Amsterdam, The Netherlands; Brussels, Belgium; Vienna, Austria; Madrid, Spain; Rome, Italy; Bern, Switzerland; Stockholm, Sweden; Oslo, Norway
+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	
+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 Configuration Commands

Table 19 describes SNTP configuration commands.

Table 19: SNTP Configuration Commands

Command	Description
config sntp-client [primary secondary] server [<ipaddress> <host_name>]	Configures an NTP server for the switch to obtain time information. 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 second server.

Table 19: SNTP Configuration Commands (continued)

Command	Description
<code>config sntp-client update-interval <seconds></code>	Configures the interval between polling for time information from SNTP servers. The default setting is 64 seconds.
<code>disable sntp-client</code>	Disables SNTP client functions.
<code>enable sntp-client</code>	Enables Simple Network Time Protocol (SNTP) client functions.
<code>show sntp-client</code>	Displays configuration and statistics for the SNTP client.

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:

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

6

Configuring Ports on a Switch

This chapter describes the following topics:

- Enabling and Disabling Switch Ports on page 71
- Load Sharing on the Switch on page 73
- Switch Port-Mirroring on page 75
- Extreme Discovery Protocol on page 76

Enabling and Disabling Switch Ports

By default, all ports are enabled. To enable or disable one or more ports, use the following command:

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

For example, to disable ports 3, 5, and 12 through 15 on a Summit24e3 switch, 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

By default, the switch is configured to use autonegotiation to determine the port speed and duplex setting for each port. You can manually configure the duplex setting and the speed of 10/100 Mbps ports.

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 ports are statically set to 1 Gbps, and their speed cannot be modified.

All ports on a stand-alone switch can be configured for half-duplex or full-duplex operation. By default, the 10/100 Mbps ports autonegotiate the duplex setting.

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

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

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

```
config ports <portlist> auto on
```

Flow control is supported only on Gigabit Ethernet ports. It is enabled or disabled as part of autonegotiation. If autonegotiation is set to off, flow control is disabled. When autonegotiation is turned on, flow control is enabled.

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.

The following example turns autonegotiation off for port 4 (a Gigabit Ethernet port) on a stand-alone switch:

```
config ports 4 auto off duplex full
```

Switch Port Commands

Table 20 describes the switch port commands.

Table 20: Switch Port Commands

Command	Description
config ports <portlist> auto off {speed [10 100 1000]} duplex [half full]	Changes the configuration of a group of ports. Specify the following: <code>auto off</code> — The port will not autonegotiate the settings. <code>speed</code> — The speed of the port. <code>duplex</code> — The duplex setting (half- or full-duplex).
config ports <portlist> auto on	Enables autonegotiation for the particular port type; 802.3u for 10/100 Mbps ports or 802.3z for Gigabit Ethernet ports.
config ports <portlist> display-string <string>	Configures a user-defined string for a port. The string is displayed in certain <code>show</code> commands (for example, <code>show port all info</code>). The string can be up to 16 characters.
config sharing address-based [mac_source mac_destination mac_source_destination ip_source ip_destination ip_source_destination]	Configures the part of the packet examined by the switch when selecting the egress port for transmitting load-sharing data. This feature is available using the address-based load-sharing algorithm, only.
disable ports <portlist>	Disables a port. Even when disabled, the link is available for diagnostic purposes.
disable sharing <port>	Disables a load-sharing group of ports.
enable ports <portlist>	Enables a port.
enable sharing <port> grouping <portlist> {address-based}	Defines a load-sharing group of ports. The ports specified in <portlist> are grouped to the master port. The optional load-sharing algorithm, address-based, uses addressing information as criteria for egress port selection.

Table 20: Switch Port Commands (continued)

Command	Description
restart ports <portlist>	Resets autonegotiation for one or more ports by resetting the physical link.
show ports {<portlist>} collisions	Displays real-time collision statistics.
show ports {<portlist>} configuration	Displays the port configuration.
show ports {<portlist>} info {detail}	Displays detailed system-related information.
show ports {<portlist>} packet	Displays a histogram of packet statistics.
show ports {<portlist>} rxerrors	Displays real-time receive error statistics.
show ports {<portlist>} stats	Displays real-time port statistics.
show ports {<portlist>} txerrors	Displays real-time transmit error statistics.
show ports {<portlist>} utilization	Displays real-time port utilization information. Use the [Spacebar] to toggle between packet, byte, and bandwidth utilization information.
show sharing address-based	Displays the address-based load sharing configuration.
unconfig ports <portlist> display-string <string>	Clears the user-defined display string from a port.

Load Sharing on the Switch

Load sharing with switches allows you to increase bandwidth and resiliency by using a group of ports to carry traffic in parallel between switches. The 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 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.

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 Algorithms

Load-sharing algorithms allow you to select the distribution technique used by the load-sharing group to determine the output port selection. Algorithm selection is not intended for use in predictive traffic engineering.

You can configure the address-based load-sharing algorithm on the Summit24e3 switch.

The address-based load-sharing 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:

- For Layer 2 load sharing, the switch uses the MAC source address, MAC destination address, IP source address, and IP destination address.
- For Layer 3 load sharing, the switch uses the IP destination address.

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

```
config sharing address-based [mac_source | mac_destination | mac_source_destination |
ip_source | ip_destination | ip_source_destination]
```

where:

- `mac_source` — Indicates that the switch should examine the MAC source address.
- `mac_destination` — Indicates that the switch should examine the MAC destination address.
- `mac_source_destination` — Indicates that the switch should examine the MAC source and destination address.
- `ip_source` — Indicates that the switch should examine the IP source address.
- `ip_source_destination` — Indicates that the switch should examine the IP source address and destination address.
- `ip_destination` — Indicates that the switch should examine the IP destination address.

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.

The following rules apply to the Summit24e3 switch:

- Ports on the switch must be of the same port type. For example, if you use 100 Mbps ports, all ports on the switch must be 100 Mbps ports.
- Ports on the switch are divided into a maximum of eight groups.

- Port-based and round-robin load sharing algorithms do not apply.

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> {address-based}
disable sharing <port>
```

Load-Sharing Example

This section provides an example of how to define load-sharing on a Summit24e3 switch.

Load-Sharing on a Summit24e3 Switch

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.



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.

Verifying the Load-Sharing Configuration

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

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.

The traffic filter is defined by the following criteria:

- **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 always transmits tagged frames. The default port tag will be added to any untagged packets as they are mirrored. 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).

**NOTE**

For optimum performance, mirror three or fewer ports at any given time.

Port-Mirroring Commands

Switch port-mirroring commands are described in Table 21.

Table 21: Switch Port-Mirroring Configuration Commands

Command	Description
config mirroring add ports <portlist>	Adds a single mirroring filter definition. Up to eight mirroring definitions can be added.
config mirroring delete ports <portlist>	Deletes a particular mirroring filter definition.
disable mirroring	Disables port-mirroring.
enable mirroring to <port> tagged	Dedicates a port to be the mirror output port.
show mirroring	Displays the port-mirroring configuration.

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
config mirroring add port 1
```

Extreme Discovery Protocol

The Extreme Discovery Protocol (EDP) is used to gather information about neighbor Extreme Networks switches. EDP is used to by the switches to exchange topology information. 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 Commands

Table 22 lists EDP commands.

Table 22: EDP Commands

Command	Description
disable edp ports <portlist>	Disables the EDP on one or more ports.
enable edp ports <portlist>	Enables the generation and processing of EDP messages on one or more ports. The default setting is enabled.
show edp	Displays EDP information.



Virtual LANs (VLANs)

This chapter describes the following topics:

- Overview of Virtual LANs on page 79
- Types of VLANs on page 80
- VLAN Names on page 84
- Configuring VLANs on the Switch on page 85
- Displaying VLAN Settings on page 86

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
- 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. A port can be a member of only one port-based VLAN. The Summit24e3 switch supports L2 port-based VLANs.

For example, on the Summit24e3 switch in Figure 7, ports 1 through 8, and port 26 are part of VLAN *Sales*; ports 9 through 16, and port 25 are part of VLAN *Finance*; and ports 17 through 24 are part of VLAN *Marketing*.

Figure 7: Example of a port-based VLAN on the Summit24e3 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 8 illustrates a single VLAN that spans a BlackDiamond switch and a Summit24e3 switch. All ports on the BlackDiamond switch belong to VLAN *Sales*. Ports 1 through 24, and port 26 on the Summit 24e3 switch also belong to VLAN *Sales*. The two switches are connected using slot 8, port 4 on system 1 (the BlackDiamond switch), and port 26 on system 2 (the Summit24e3 switch).

Figure 8: Single port-based VLAN spanning two switches

SH_004

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 9 illustrates two VLANs spanning two switches. On system 1, ports 1 through 8, and port 26 are part of VLAN *Accounting*; ports 17 through 24, and port 25 are part of VLAN *Engineering*. On system 2, all ports on slot 1 are part of VLAN *Accounting*; all ports on slot 8 are part of VLAN *Engineering*.

Figure 9: Two port-based VLANs spanning two switches

SH_005

VLAN *Accounting* spans system 1 and system 2 by way of a connection between system 1, port 26 and system 2, slot 1, port 6. VLAN *Engineering* spans system 1 and system 2 by way of a connection between system 1, port 25, and system 2, 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*. The Summit24e3 switch supports L2 tagged VLANs.



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 9. 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 10 illustrates the physical view of a network that uses tagged and untagged traffic.



Figure 10: Physical diagram of tagged and untagged traffic

Figure 11 is a logical diagram of the same network.

Figure 11: Logical diagram of tagged and untagged traffic

In Figure 10 and Figure 11:

- 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 16 on system 1 has a NIC that supports 802.1Q tagging.
- The server connected to port 16 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.



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 does not begin with an alphabetical character, or that contains a space, comma, or other special character.

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.



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:

```
config 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.

- 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 Commands

Table 23 describes the commands used to configure a VLAN.

Table 23: VLAN Configuration Commands

Command	Description
config vlan <name> add port <portlist> {tagged untagged} {nobroadcast}	Adds one or more ports to a VLAN. You can specify tagged port(s), untagged port(s). Specify <i>nobroadcast</i> to prevent the switch from forwarding broadcast, multicast, and unknown unicast traffic. By default, ports are untagged.
config vlan <name> delete port <portlist> {tagged untagged} {nobroadcast}	Deletes one or more ports from a VLAN.
config vlan <name> ipaddress <ipaddress> {<mask>}	Assigns an IP address and an optional mask to the VLAN.
config vlan <name> tag <vlanid>	Assigns a numerical VLANid. The valid range is from 2 to 4094 (1 is used by the default VLAN).

Table 23: VLAN Configuration Commands (continued)

Command	Description
<code>config vlan <old_name> name <new_name></code>	Renames a previously configured VLAN.
<code>create vlan <name></code>	Creates a named VLAN.
<code>delete vlan <name></code>	Removes a VLAN.
<code>unconfig ports <portlist> monitor vlan <name></code>	Removes port-based VLAN monitoring.
<code>unconfig vlan <name> ipaddress</code>	Resets the IP address of the VLAN.

VLAN Configuration Examples

The following Summit24e3 switch 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
config video tag 1000
config video add port 4-8 tagged
```

The following Summit24e3 switch 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
config sales tag 120
config sales add port 1-3 tagged
config sales add port 4,7
```

Displaying VLAN Settings

To display VLAN settings, use the following command:

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

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

- Name
- VLANid
- How the VLAN was created
- IP address
- STPD 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.

8

Forwarding Database (FDB)

This chapter describes the following topics:

- Overview of the FDB on page 87
- Configuring FDB Entries on page 88
- Displaying FDB Entries on page 90

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 on which it was received, and an identifier for the VLAN to which the device belongs. Frames destined for devices that are not in the FDB are flooded to all members of the VLAN.

FDB Entry Types

The Summit24e3 supports up to 8,191 layer 2 FDB entries and 2,047 layer 3 FDB entries. The following are four types of entries in the FDB:

- **Dynamic entries** — Initially, all entries in the database are dynamic. Entries in the database are removed (aged-out) if, after a period of time (aging time), the device has not transmitted. 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. Dynamic entries are deleted from the database if the switch is reset or a power off/on cycle occurs. For more information about setting the aging time, refer to “Configuring FDB Entries” later in this chapter.
- **Nonaging entries** — If the aging time is set to zero, all aging entries in the database are defined as static, nonaging 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. The system administrator must make entries permanent. A permanent entry can either be a unicast or multicast MAC address. All entries entered by way of the command-line

interface are stored as permanent. The Summit24e3 switches support a maximum of 64 permanent entries.

Once created, permanent entries stay the same as when they were created. For example, the permanent entry store is not 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).
- **Blackhole entries** — A blackhole entry configures the switch to discard packets with a specified MAC destination address. Blackhole entries are useful as a security measure or in special circumstances where a specific destination address must be discarded. 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.

How FDB Entries Get Added

Entries are added into the FDB in the following two ways:

- The switch can learn entries. 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.
- You can enter and update entries using a MIB browser, an SNMP Network Manager, or the command-line interface (CLI).

Associating a QoS Profile with an FDB Entry

You can associate a QoS profile with a MAC address (and VLAN) of a device that will be dynamically learned. The FDB treats the entry like a dynamic entry (it is learned, it can be aged out of the database, and so on). The switch applies the QoS profile as soon as the FDB entry is learned.



NOTE

For more information on QoS, refer to Chapter 12.

Configuring FDB Entries

To configure entries in the FDB, use the commands listed in Table 24.

Table 24: FDB Configuration Commands

Command	Description
<code>clear fdb [{<mac_address> vlan <name> ports <portlist>}]</code>	Clears dynamic FDB entries that match the filter. When no options are specified, the command clears all FDB entries.
<code>config fdb agingtime <number></code>	Configures the FDB aging time. The range is 15 through 1,000,000 seconds. The default value is 300 seconds. A value of 0 indicates that the entry should never be aged out.
<code>create fdbentry <mac_address> vlan <name> ports [<portlist> all] {{qosprofile <qosprofile> {ingress-qosprofile <qosprofile>}} {ingress-qosprofile <qosprofile> {qosprofile <qosprofile>}}</code>	Creates a permanent static FDB entry. Specify the following: <code>mac_address</code> — Device MAC address, using colon separated bytes. <code>name</code> — VLAN associated with MAC address. <code>portlist</code> — Port numbers associated with MAC address. <code>qosprofile</code> — QoS profile associated with destination MAC address of the egress port. <code>ingress-qosprofile</code> — QoS profile associated with the source MAC address of the ingress port. If more than one port number is associated with a permanent MAC entry, packets are multicast to the multiple destinations.
<code>create fdbentry <mac_address> vlan <name> dynamic {{qosprofile <qosprofile> {ingress-qosprofile <qosprofile>}} {ingress-qosprofile <qosprofile> {qosprofile <qosprofile>}}</code>	Creates a permanent dynamic FDB entry. Assigns a packet with the specified MAC address and VLAN to a specific QoS profile. If you only specify 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.
<code>create fdbentry <mac_address> vlan <name> blackhole {source-mac dest-mac both}</code>	Creates a blackhole FDB entry. Specify: <code>source-mac</code> — The blackhole MAC address matches the ingress source MAC address. <code>dest-mac</code> — The blackhole MAC address matches the egress destination MAC address. <code>both</code> — The blackhole MAC address matches the ingress source MAC address or the egress destination MAC address.
<code>delete fdbentry {<mac_address> vlan <name> all}</code>	Deletes one or all permanent FDB entries.
<code>disable learning port <portlist></code>	Disables MAC address learning on one or more ports for security purposes. 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. The default setting is enabled.
<code>enable learning port <portlist></code>	Enables MAC address learning on one or more ports.

FDB Configuration Examples

The following example adds a permanent entry to the FDB:

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

The permanent entry has the following characteristics:

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

This example associates the QoS profile *qp2* with a dynamic entry 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 00A023123456.
- VLAN name is *net34*.
- The entry will be learned dynamically.
- QoS profile *qp2* will be applied when the entry is learned.

Displaying FDB Entries

To display FDB entries, use the following command:

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

where the following is true:

- `mac_address` — Displays the entry for a particular MAC address.
- `vlan <name>` — Displays the entries for a VLAN.
- `ports <portlist>` — Displays the entries for a slot and port combination.
- `permanent` — Displays all permanent entries, including the ingress and egress QoS profiles.

With no options, the command displays all FDB entries.



Access Policies

This chapter describes the following topics:

- Overview of Access Policies on page 91
- Using Access Control Lists on page 92
- Using Routing Access Policies on page 104
- Making Changes to a Routing Access Policy on page 108
- Removing a Routing Access Policy on page 108
- Routing Access Policy Commands on page 109

Overview of Access Policies

Access policies are a generalized category of features that impact forwarding and route forwarding decisions. Access policies are used primarily for security and quality of service (QoS) purposes.

The three categories of access policies are:

- Access control lists
- Rate limits
- Routing access policies

Access Control 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 DiffServe code point. Using access lists has no impact on switch performance.

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.

Routing Access Policies

Routing access policies are used to control the advertisement or recognition of routing protocols, such as RIP or OSPF. Routing access policies can be used to 'hide' entire networks, or to trust only specific sources for routes or ranges of routes. The capabilities of routing access policies are specific to the type of routing protocol involved, but are sometimes more efficient and easier to implement than access lists.

Using Access Control Lists

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 following sections describe how to use access control lists.

Access Masks

There are between twelve and fourteen access masks available in the Summit24e3, 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).

An access mask consists of a combination of the following thirteen fields:

- Ethernet destination MAC address
- Ethernet source MAC address
- VLANid
- IP Type of Service (TOS) or DiffServ code point
- Ethertype
- IP protocol
- IP destination address and netmask
- Layer 4 destination port
- IP source address and netmask
- Layer 4 source port, or ICMP type and/or ICMP code
- TCP session initiation bits (permit-established keyword)
- Egress port
- Ingress ports

An access mask can also have an optional, unique precedence number associated with it.

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.

For packets that match a particular access control list, you can specify the following actions:

- Drop
Drop the packets. 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 95).
- Permit
Forward the packet. You can send the packet to a particular QoS profile, and modify the packet’s 802.1p value and / or DiffServe code point.

Rate Limits

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.



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.

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 DiffServe 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 DiffServe code point.

The allowable rate limit values for the 100BT ports are 1, 2, 3, 4 ... 100 Mbps, and for the Gigabit ports are 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.

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 DiffServe code point.

Access Mask Precedence Numbers

The access mask precedence number is optional, and determines the order in which each rule is examined by the switch. 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*. However, 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 Summit24e3 port used as a management port.

The following example shows an access control list that is used to specify an default rule to explicitly deny all traffic:

```
create access-mask ingress_mask ports precedence 25000
create access-list DenyAll ingress_mask ports 2-26 deny
```

Once the default behavior of the access control list is established, you can create additional entries using precedence numbers.

The following access control list example shows an access control list that will forward traffic from the 10.1.2.x subnet even while the above default rule is in place:

```
create access-mask ip_src_mask source-ip/24 precedence 1000
create access-list TenOneTwo ip_src_mask source-ip 10.1.2.0/24 permit
```

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.



NOTE

For an example of using the permit-established keyword, refer to “Using the Permit-Established Keyword” on page 100.

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 (see Table 25 for the full command syntax). 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 | ports <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 Commands

Table 25 describes the commands used to configure access control lists.

Table 25: Access Control List Configuration Commands

Command	Description
<pre> 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 <protocol_num>}} {dest-ip <dest_IP>/<mask length>} {dest-L4port <dest_port>} {source-ip <src_IP>/<mask length>} {source-L4port <src_port> {icmp-type <icmp_type>} {icmp-code <icmp_code>}} {egressport <port>} {ports <portlist>} [permit {qosprofile <qosprofile>} {set code-point <code_point>} {set dot1p <dot1p_value>} permit-established deny] </pre>	Creates an access list. The list is applied to all ingress packets. Options include: <name> — Specifies the access control list name. The access list name can be between 1 and 31 characters. access-mask — Specifies the associated access mask. Any field specified in the access mask must have a corresponding value specified in the access list. dest-mac — Specifies the destination MAC address. source-mac — Specifies the source MAC address. vlan — Specifies the VLANid. ethertype — Specify IP, ARP, or the hex value to match. tos — Specifies the IP precedence value. code-point — Specifies the DiffServ code point value. ipprotocol — Specify an IP protocol, or the protocol number dest-ip — Specifies an IP destination address and subnet mask. A mask length of 32 indicates a host entry. dest-L4port — Specify the destination port. source-ip — Specifies an IP source address and subnet mask. source-L4port — Specify the source port. icmp-type — Specify the ICMP type. icmp-code — Specify the ICMP code. egressport — Specify the egress port ports — Specifies the ingress port(s) on which this rule is applied. permit — Specifies the packets that match the access list description are permitted to be forward by this switch. An optional QoS profile can be assigned to the access list, so that the switch can prioritize packets accordingly. set — Modify the DiffServ code point and/or the 802.1p value for matching packets. permit-established — Specifies a uni-directional session establishment is denied. deny — Specifies the packets that match the access list description are filtered (dropped) by the switch.

Table 25: Access Control List Configuration Commands (continued)

Command	Description
<pre>create access-mask <access-mask name> {dest-mac} {source-mac} {vlan } {ethertype} {tos code-point} {ipprotocol} {dest-ip /<mask length>} {dest-L4port} {source-ip /<mask length>} {source-L4port {icmp-type} {icmp-code}} {permit-established} {egressport} {ports} {precedence <number>}</pre>	Creates an access mask. The mask specifies which packet fields to examine. Options include: <code><access-mask name></code> — Specifies the access mask name. The access mask name can be between 1 and 31 characters. <code>dest-mac</code> — Specifies the destination MAC address field. <code>source-mac</code> — Specifies the source MAC address field. <code>vlan</code> — Specifies the VLANid field. <code>ethertype</code> — Specifies the Ethertype field. <code>tos</code> — Specifies the IP precedence field. <code>code-point</code> — Specifies the DiffServ code point field. <code>ipprotocol</code> — Specifies the IP protocol field. <code>dest-ip</code> — Specifies the IP destination field and subnet mask. You must supply the subnet mask. <code>dest-L4port</code> — Specifies the destination port field. <code>source-ip</code> — Specifies the IP source address field and subnet mask. You must supply the subnet mask. <code>source-L4port</code> — Specifies the source port field. <code>icmp-type</code> — Specify the ICMP type field. <code>icmp-code</code> — Specify the ICMP code field. <code>permit-established</code> — Specifies the TCP SYN/ACK bit fields. <code>egressport</code> — Specify the egress port <code>ports</code> — Specifies the ingress port(s) on which this rule is applied. <code>precedence</code> — Specifies the access mask precedence number. The range is 1 to 25,600.

Table 25: Access Control List Configuration Commands (continued)

Command	Description
<pre> create rate-limit <rule_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]<protocol_num>}} {dest-ip <dest_IP>/<mask length>} {dest-L4port <dest_port>} {source-ip <src_IP>/<mask length>} {source-L4port <src_port> {icmp-type <icmp_type>} {icmp-code <icmp_code>}} {egressport <port>} {port <port number>} permit {qosprofile <qosprofile>} {set code-point <code_point>} {set dot1p <dot1p_value>} limit <rate_in_Mbps> {exceed-action [drop set code-point <code_point>} </pre>	Creates a rate limit. The rule is applied to all ingress packets. Options include: • <rule_name> — Specifies the rate limit name. The name can be between 1 and 31 characters. • access-mask — Specifies the associated access mask. Any field specified in the access mask must have a corresponding value specified in the rate limit. • dest-mac — Specifies the destination MAC address. • source-mac — Specifies the source MAC address. • vlan — Specifies the VLANid. • ethertype — Specify IP, ARP, or the hex value to match. • tos — Specifies the IP precedence value. • code-point — Specifies the DiffServ code point value. • ipprotocol — Specify an IP protocol, or the protocol number • dest-ip — Specifies the IP destination address and subnet mask. A mask length of 32 indicates a host entry. • dest-L4port — Specify the destination port. • source-ip — Specifies the IP source address and subnet mask. • source-L4port — Specify the source port. • icmp-type — Specify the ICMP type. • icmp-code — Specify the ICMP code. • egressport — Specify the egress port • port — Specifies the ingress port to which this rule is applied. • permit — Specifies the packets that match the access list description are permitted to be forward by this switch. An optional QoS profile can be assigned to the access list, so that the switch can prioritize packets accordingly. • set — Modify the DiffServ code point or the 802.1p value for matching, forwarded, packets. • limit — Specifies the rate limit • <rate_in_Mbps> — The rate limit. Allowed values are 1-100 Mbps for 100BT ports, 8, 16, 24, 32... 1000 for the Gigabit ports • exceed-action — Action to take for matching packets that exceed the rate.
<pre> delete access-list <name> </pre>	Deletes an access list.

Table 25: Access Control List Configuration Commands (continued)

Command	Description
delete access-mask <name>	Deletes an access mask. Any access lists or rate limits that reference this mask must first be deleted.
delete rate-limit <name>	Deletes a rate limit.
show access-list {<name> ports <portlist>}	Displays access-list information.
show access-mask {<name>}	Displays access-list information.
show rate-limit {<name> ports <portlist>}	Displays access-list information.

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 12, 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.
- IPForwarding is enabled.

Figure 12: 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 13 illustrates the outcome of the access control list.

Figure 13: 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 qpl
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 qpl
```

Figure 14 illustrates the outcome of this access list.

Figure 14: 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 15 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 15: 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 15. 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 16 shows the final outcome of this access list.


Figure 16: 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 deny ping icmp_mask ipprotocol icmp icmp-type 8 icmp-code 0 deny
```

The output for this access list is shown in Figure 17.

Figure 17: 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
```

Using Routing Access Policies

To use routing access policies, 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 are IP addresses and subnet masks
- 4 Apply the access profile.

Creating an Access Profile

The first thing to do when using routing access policies is to create an *access profile*. An access profile has a unique name and contains a list of IP addresses and associated subnet masks.

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
```

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:

```
config access-profile <access_profile> mode [permit | deny | none]
```

Adding an Access Profile Entry

Next, configure the access profile, using the following command:

```
config access-profile <access_profile> add {<seq_number>} {permit | deny} [ipaddress  
<ipaddress> <mask> {exact}]
```

The following sections describe the `config 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 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.

Deleting an Access Profile Entry

To delete an access profile entry, use the following command:

```
config 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 Access Policies 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:

```
config rip 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:

```
config rip 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:

```
config rip vlan [<name> | all] export-filter [<access_profile> | none]
```

Examples

In the example shown in Figure 18, 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 18: 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
config access-profile nointernet mode deny
config access-profile nointernet add 10.0.0.10/32
config 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
config access-profile nosales mode deny
config access-profile nosales add 10.2.1.0/24
```

```
config rip vlan backbone import-filter nosales
```

This configuration results in the switch having no route back to the VLAN Sales.

Routing Access Policies for OSPF

Because OSPF is a link-state protocol, the access policies associated with OSPF are different in nature than those associated with RIP. Access policies 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:

```
config ospf area <area_id> 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:

```
config ospf area <area_id> 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:


```
config 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:

```
config ospf direct-filter [<access_profile> | none]
```

Example

Figure 19 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 19: OSPF access policy example

EW_002

To configure the switch labeled Internet, the commands would be as follows:

```
create access-profile okinternet ipaddress
config access-profile okinternet mode permit
config access-profile okinternet add 192.1.1.0/24
config ospf asbr-filter okinternet
```

Making Changes to a Routing Access Policy

You can change the routing access policy by changing the associated access profile. However, the propagation of the change depends on the protocol and policy involved. Propagation of changes applied to RIP access policies depends on the protocol timer to age-out entries.



NOTE

Changes to profiles applied to OSPF typically require rebooting the switch, or disabling and re-enabling OSPF on the switch.

Removing a Routing Access Policy

To remove a routing access policy, you must remove the access profile from the routing protocol or VLAN. All the commands that apply an access profile to form an access policy also have the option of choosing none as the access profile. Using the none option removes any access profile of that particular type from the protocol or VLAN, and, therefore, removes the access policy.

Routing Access Policy Commands

Table 26 describes the commands used to configure routing access policies.

Table 26: Routing Access Policy Configuration Commands

Command	Description
<pre>config access-profile <access_profile> add {<seq_number>} {permit deny} [ipaddress <ipaddress> <mask> {exact}]</pre>	Adds an entry to the access profile. The explicit sequence number, and permit or deny attribute should be specified if the access profile mode is none. Specify one of the following: • <seq-number> — The order of the entry within the access profile. If no sequence number is specified, the new entry is added to the end of the access-profile and is automatically assigned a value of 5 more than the sequence number of the last entry. • permit deny — Per-entry permit or deny specification. The per-entry attribute only takes effect if the access-profile mode is none. Otherwise, the overall access profile type takes precedence. • <ipaddress> <mask> — An IP address and mask. If the attribute "exact" is specified for an entry, then a exact match with address and mask is performed, subnets within the address range do not match entry against entry. •
<pre>config access-profile <access_profile> delete <seq_number></pre>	Deletes an access profile entry using the sequence number.
<pre>config access-profile <access_profile> mode [permit deny none]</pre>	Configures the access profile to be one of the following: • permit — Allows the addresses that match the access profile description. • deny — Denies the addresses that match the access profile description. • none — Permits and denies access on a per-entry basis. Each entry must be added to the profile as either type permit or deny. The default setting is permit.
<pre>config ospf area <area_id> external-filter [<access_profile> none]</pre>	Configures the router to use the access policy to determine which external routes are allowed to be exported into the area. This router must be an ABR.
<pre>config ospf area <area_id> interarea-filter [<access_profile> none]</pre>	Configures the router to use the access policy to determine which inter-area routes are allowed to be exported into the area. This router must be an ABR.

Table 26: Routing Access Policy Configuration Commands (continued)

Command	Description
config ospf asbr-filter [<access_profile> none]	Configures the router to use the access policy to limit the routes that are advertised into OSPF for the switch as a whole for switches configured to support RIP and static route re-distribution into OSPF.
config ospf direct-filter [<access_profile> none]	Configures the router to use the access policy to limit the routes that are advertised into OSPF for the switch as a whole for switches configured to support direct route re-distribution into OSPF.
config rip vlan [<name> all] export-filter [<access-profile> none]	Configures RIP to suppress certain routes when performing route advertisements.
config rip vlan [<name> all] import-filter [<access_profile> none]	Configures RIP to ignore certain routes received from its neighbor.
config rip vlan [<name> all] trusted-gateway [<access_profile> none]	Configures RIP to use the access list to determine which RIP neighbor to receive (or reject) the routes.
create access-profile <access_profile> type [ipaddress]	Creates an access profile. Once the access profile is created, one or more addresses can be added to it, and the profile can be used to control a specific routing protocol. Specify the following: • ipaddress — A list of IP address and mask pairs. •
delete access-profile <access_profile>	Deletes an access profile.
show access-profile <access_profile>	Displays access-profile related information for the switch.

10

Network Address Translation (NAT)


This chapter covers the following topics:

- Overview on page 111
- Internet IP Addressing on page 112
- Configuring VLANs for NAT on page 112
- Configuring NAT on page 114
- Displaying NAT Settings on page 117
- Disabling NAT on page 117

Overview

NAT is a feature that allows one set of IP addresses, typically private IP addresses, to be converted to another set of IP addresses, typically public Internet IP addresses. This conversion is done transparently by having a NAT device rewrite the source IP address and Layer 4 port of the packets.

Figure 20: NAT Overview



You can configure NAT to conserve IP address space by mapping a large number of inside (private) addresses to a much smaller number of outside (public) addresses.

In implementing NAT, you must configure at least two separate VLANs involved. One VLAN is configured as inside, and corresponds to the private IP addresses you would like to translate into other

IP addresses. The other type of VLAN is configured as outside, which corresponds to the public (probably Internet) IP addresses you want the inside addresses translated to. The mappings between inside and outside IP addresses are done via rules that specify the IP subnets involved and the algorithms used to translate the addresses.

**NOTE**

The NAT modes in ExtremeWare support translating traffic initiating only from inside addresses.

NAT rules are associated with a single outside VLAN. Multiple rules per outside VLAN are allowed. The rules take effect in the order they are displayed using the `show` command. Any number of inside VLANs can use a single outside VLAN, assuming that you have created proper rules. Similarly, a single inside VLAN can use any number of different outside VLANs, assuming that the rules and routing are set up properly.

Both TCP and UDP have Layer 4 port numbers ranging from 1 to 65535. These Layer 4 ports, in combination with the IP addresses, form a unique identifier which allows hosts (as well as the NAT switch) to distinguish between separate conversations. NAT operates by replacing the inside IP packet's source IP and Layer 4 port with an outside IP and Layer 4 port. The NAT switch maintains a connection table to map the return packets on the outside VLAN back into their corresponding inside sessions.

Internet IP Addressing

When implementing NAT in an Internet environment, it is strongly recommended that you use one of the reserved private IP address ranges for your inside IP addresses. These ranges have been reserved specifically for networks not directly attached to the Internet. Using IP addresses within these ranges prevents addressing conflicts with public Internet sites to which you want to connect. The ranges are as follows:

- 10.0.0.0/8—Reserved Class A private address space
- 172.16.0.0/12—Reserved Class B private address space
- 192.168.0.0/16—Reserved Class C private address space

Configuring VLANs for NAT

You must configure each VLAN participating in NAT as either an inside or outside VLAN. To configure a VLAN as an inside or outside VLAN, use the following command:

```
config nat vlan <name> [inside | outside | none]
```

When a VLAN is configured to be `inside`, traffic from that VLAN destined for an `outside` VLAN is translated only if it has a matching NAT rule. Any unmatched traffic will be routed normally and not be translated. Because all traffic destined for an `outside` VLAN runs through the central processing unit (CPU), it cannot run at line-rate.

When a VLAN is configured to be `outside`, it routes all traffic destined for `inside` VLANs. Because the routed traffic runs through the CPU, it cannot run at line-rate.

When a VLAN is configured to be `none`, all NAT functions are disabled and the VLAN operates normally.

NAT Modes

There are 4 different modes used to determine how the outside IP addresses and Layer 4 ports are assigned.

- Static mapping
- Dynamic mapping
- Port-mapping
- Auto-constraining

Static Mapping

When static mapping is used, each inside IP address uses a single outside IP address. The Layer 4 ports are not changed, only the IP address is rewritten. Because this mode requires a 1-to-1 mapping of internal to external addresses, it does not make efficient use of the external address space. However, it is useful when you have a small number of hosts that need to have their IP addresses rewritten without conflicting with other hosts. Because this mode does not rely on Layer 4 ports, ICMP traffic is translated and allowed to pass.

Dynamic Mapping

Dynamic mapping is similar to static mapping in that the Layer 4 ports are not rewritten during translation. Dynamic mapping is different in that the number of inside hosts can be greater than the number of outside hosts. The outside IP addresses are allocated on a first-come, first-serve basis to the inside IP addresses. When the last session for a specific inside IP address closes, that outside IP address can be used by other hosts. Since this mode does not rely on Layer 4 ports, ICMP traffic is translated and allowed to pass.

Port-mapping

Port-mapping gives you the most efficient use of the external address space. As each new connection is initiated from the inside, the NAT device picks the next available source Layer 4 port on the first available outside IP address. When all ports on a given IP address are in use, the NAT device uses ports off of the next outside IP address. Some systems reserve certain port ranges for specific types of traffic, so it is possible to map specific source Layer 4 port ranges on the inside to specific outside source ranges. However, this may cause a small performance penalty. In this case, you would need to make several rules using the same inside and outside IP addresses, one for each Layer 4 port range. ICMP traffic is not translated in this mode. You must add a dynamic NAT rule for the same IP address range to allow for ICMP traffic.

Auto-constraining

The auto-constraining algorithm for port-mapping limits the number of outside Layer 4 ports a single inside host can use simultaneously. The limitation is based on the ratio of inside to outside IP addresses. The outside IP address and Layer 4 port space is evenly distributed to all possible inside hosts. This guarantees that no single inside host can prevent other traffic from flowing through the NAT device. Because of the large number of simultaneous requests that can be made from a web browser, it is not recommended that this mode be used when a large number of inside hosts are being translated to a small number of outside IP addresses. ICMP traffic is not translated in this mode. You must add a dynamic NAT rule for the same IP address range to allow for ICMP traffic.

Configuring NAT

The behavior of NAT is determined by the rules you create to translate the IP addresses. You must attach each rule to a specific VLAN. All rules are processed in order. The options specified on the NAT rule determine the algorithm used to translate the inside IP addresses to the outside IP addresses. For outgoing (inside to outside) packets, the first rule to match is processed. All following rules are ignored. All return packets must arrive on the same outside VLAN on which the session went out. For most configurations, make sure that the outside IP addresses specified in the rule are part of the outside VLAN's subnet range, so that the switch can proxy the address resolution protocol (ARP) for those addresses.

To enable NAT functionality, use the following command:

```
enable nat
```

Configuring NAT Rules

To configure NAT rules, use the commands listed in Table 27.

Table 27: NAT Configuration Commands

Command	Description
<pre>config nat add vlan <outside_vlan> map source [any <ipaddress> [/<bits> <netmask>]] {l4-port [any <number> {- <number>}]} {destination <ipaddress>/<mask> {l4-port [any <number> {- <number>}]} to <ipaddress> [/<mask> <netmask> - <ipaddress>] {[tcp udp both] [portmap {<min> - <max>} auto-constrain]}</pre>	Adds a NAT translation rule that translates private IP addresses to public IP addresses on the outside VLAN. The first IP address specifies private side IP addresses and the second IP address specifies the public side IP address. Use portmap to specify port translations and specify either TCP or UDP port translation, or both. The range of <code>number</code> is 1 to 65535. The default setting for <code>min</code> is 1024. The default setting for <code>max</code> is 65535.
<pre>config nat delete vlan <outside_vlan> map source [any <ipaddress> [/<bits> <netmask>]] {l4-port [any <number> {- <number>}]} {destination <ipaddress>/<mask> {l4-port [any <number> {- <number>}]} to <ipaddress> [/<mask> <netmask> - <ipaddress>] {[tcp udp both] [portmap {<min> - <max>} auto-constrain]}</pre>	Deletes a NAT translation rule.

Creating NAT Rules

This section describes how to configure the various types of NAT (static, dynamic, portmap, and auto-constrain). In the examples in this section, advanced port and destination matching options have been removed. For information on how to use some of the more advanced rule matching features, refer to “Advanced Rule Matching” on page 116.

Creating Static and Dynamic NAT Rules

To create static or dynamic NAT rules, use this command:

```
config nat [add | delete] vlan <outside_vlan> map source [any | <ipaddress>
[/<bits>| <netmask>]] to <ipaddress> [/<mask> | <netmask> | - <ipaddress>]
```

This is the simplest NAT rule. You specify the outside vlan name, and a subnet of inside IP addresses, which get translated to the outside IP address using the specified mode (static in this case). For the outside IP addresses, you can either specify an IP address and netmask or a starting and ending IP range to determine the IP addresses the switch will translate the inside IP addresses to. If the netmask for both the source and NAT addresses is /32, the switch will use static NAT translation. If the netmask for both the source and NAT addresses are not both /32, the switch will use dynamic NAT translation.

Static NAT Rule Example

```
config nat add out_vlan_1 map source 192.168.1.12/32 to 216.52.8.32/32
```

Dynamic NAT Rule Example

```
config nat add out_vlan_1 map source 192.168.1.0/24 to 216.52.8.1 - 216.52.8.31
```

Creating Portmap NAT Rules

To configure portmap NAT rules, use this command:

```
config nat [add|delete] vlan <outside_vlan> map source [any | <ipaddress> [/<bits>|
<netmask>]] to <ip> [/<mask> | <netmask> | - <ipaddress>] {[tcp |udp | both]
portmap {<min> - <max>}}
```

The addition of an L4 protocol name and the portmap keyword tells the switch to use portmap mode. Optionally, you may specify the range of L4 ports the switch chooses on the translated IP addresses, but there is a performance penalty for doing this. Remember that portmap mode will only translate TCP and/or UDP, so a dynamic NAT rule must be specified after the portmap rule in order to allow ICMP packets through without interfering with the portmapping.

Portmap NAT Rule Example

```
config nat add out_vlan_2 map source 192.168.2.0/25 to 216.52.8.32 /28 both portmap
```

Portmap Min-Max Example

```
config nat add out_vlan_2 map source 192.168.2.128/25 to 216.52.8.64/28 tcp portmap
1024 - 8192
```

Creating Auto-Constrain NAT Rules

To create auto-constrain NAT rules, use the following command:

```
config nat [add|delete] vlan <outside_vlan> map source [any | <ipaddress> [/<bits>|
<netmask>]] to <ip> [/<mask> | <netmask> | - <ipaddress>] {[tcp |udp | both]
auto-constrain}
```

This rule uses auto-constrain NAT. Remember that each inside IP address will be restricted in the number of simultaneous connections. Most installations should use portmap mode.

Auto-Constrain Example

```
config nat add out_vlan_3 map source 192.168.3.0/24 to 216.52.8.64/32 both
auto-constrain
```

Advanced Rule Matching

By default, NAT rules only match connections based on the source IP address of the outgoing packets. Using the `L4-port` and `destination` keywords, you can further limit the scope of the NAT rule so that it only applied to specific TCP/UDP Layer 4 port numbers, or specific outside destination IP addresses.



NOTE

Once a single rule is matched, no other rules are processed.

Destination Specific NAT

```
config nat [add|delete] vlan <outside_vlan> map source [any | <ipaddress> [/<bits>|
<netmask>]] {destination <ipaddress/mask> } to <ipaddress> [/<mask> | <netmask> | -
<ipaddress>]
```

The addition of the `destination` optional keyword after the source IP address and mask allows the NAT rule to be applied to only packets with a specific destination IP address.

L4-Port Specific NAT

The addition of the `L4-port` optional keyword after the source IP address and mask allows the NAT rule to be applied to only packets with a specific L4 source or destination port. If you use the `L4-port` command after the source IP/mask, the rule will only match if the port(s) specified are the source L4-ports. If you use the `L4-port` command after the destination IP/mask, the rule will only match if the port(s) specified are the destination L4-ports. Both options may be used together to further limit the rule.

Configuring Timeouts

When an inside host initiates a session, a session table entry is created. Depending on the type of traffic or the current TCP state, the table entries timeout after the configured timeout expires.

Table 28 describes the commands used to configure timeout periods.

Table 28: NAT Timeout Commands

Command	Description
<code>config nat finrst-timeout <seconds></code>	Configures the timeout for a TCP session that has been torn down or reset. The default setting is 60 seconds.
<code>config nat icmp-timeout <seconds></code>	Configures the timeout for an ICMP packet. The default setting is 3 seconds.

Table 28: NAT Timeout Commands (continued)

Command	Description
<code>config nat syn-timeout <seconds></code>	Configures the timeout for an entry with an unacknowledged TCP SYN state. The default setting is 60 seconds.
<code>config nat tcp-timeout <seconds></code>	Configures the timeout for a fully setup TCP SYN session. The default setting is 120 seconds.
<code>config nat udp-timeout <seconds></code>	Configures the timeout for an UDP session. The default setting is 120 seconds.
<code>config nat timeout <seconds></code>	Configures the timeout for any IP packet that is not TCP,UDP or ICMP. The default setting is 600 seconds.
<code>show nat timeout</code>	Displays NAT timeout settings.

Displaying NAT Settings

To display NAT rules, use the following command:

```
show nat rules {vlan <outside_vlan>}
```

This command displays the NAT rules for a specific VLAN. Rules are displayed in the order they are processed, starting with the first one.

To display NAT traffic statistics, use the following command:

```
show nat stats
```

This command displays statistics for the NAT traffic, and includes:

- The number of rules
- The number of current connections
- The number of translated packets on the inside and outside VLANs
- Information on missed translations

To display NAT connection information, use the following command:

```
show nat connections
```

This command displays the current NAT connection table, including source IP/Layer 4 port mappings from inside to outside.

Disabling NAT

To disable NAT, use the following command:

```
disable nat
```

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 on page 119
- Commands for Configuring and Monitoring EAPS on page 123


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 21).

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.

In order to use EAPS, you must enable EDP on the switch. For more information on EDP, refer to Chapter 6.


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 22), while all other nodes are designated as *transit* nodes.

Figure 21: 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 22: EAPS operation

EW_071

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.

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.

Figure 23: EAPS fault detection and protection switching

EW_072

A master node detects a ring fault in either of two ways:

- Polling response
- Trap message sent by a transit node

Polling

The master node transmits a health-check packet on the control VLAN at a user-configurable interval (see Figure 22). 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 fail-period timer and continues normal operation.

If the master node does not receive the health-check packet before the fail-period timer expires, it 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 forwarding database (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.

Trap 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 23), it immediately declares a “failed” state and performs the same steps described above; it unblocks its secondary port for access by the protected VLANs, flushes its FDB, and sends a “flush FDB” message to its associated transit nodes.

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.

Commands for Configuring and Monitoring EAPS

Table 29 lists the ExtremeWare EAPS commands. Each command is described in detail in the sections that follow.

Table 29: EAPS Commands

Command	Description
<code>config eaps <name> mode [master transit]</code>	Configures the switch as either the EAPS master node or as an EAPS transit node for the specified domain.
<code>config eaps <name> [hellotime <seconds> failtime <seconds>]</code>	Configures the values of the polling timers the master node uses for the EAPS health-check packet that is circulated around the ring for the specified EAPS domain.
<code>config eaps <name> [primary secondary] port <port number></code>	Configures a node port as the primary or secondary port for the specified EAPS domain.
<code>config eaps <name> [add delete] control vlan <name></code>	Adds the specified control VLAN to the specified EAPS domain, or deletes the specified control VLAN from the specified EAPS domain.

Table 29: EAPS Commands (continued)

Command	Description
<code>config eaps <name> [add delete] protect vlan <name></code>	Adds the specified protected VLAN to the specified EAPS domain, or deletes the specified protected VLAN from the specified EAPS domain.
<code>config eaps <old_name> name <new_name></code>	Renames an existing EAPS domain.
<code>create eaps <name></code>	Creates an EAPS domain with the specified name. Only a single domain is supported on this platform.
<code>delete eaps <name></code>	Deletes the specified EAPS domain.
<code>disable eaps</code>	Disables the EAPS function for an entire switch.
<code>disable eaps <name></code>	Disables the EAPS domain with the specified name.
<code>enable eaps</code>	Enables the EAPS function for an entire switch.
<code>enable eaps <name></code>	Enables the EAPS domain with the specified name.
<code>show eaps {<name>} [detail]</code>	Displays EAPS status information. Use the optional domain name parameter to display status information for a specific EAPS domain.
<code>unconfig eaps <name> [primary secondary] port</code>	Sets the specified port's internal configuration state to INVALID, causing the port to appear in the state Idle with a port status of Unknown when you use the <code>show eaps {<name>} detail</code> command to display port status information.

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.



NOTE

Only a single EAPS domain per switch is supported by the Summit24e3.

The following command example creates EAPS domain `eaps_1` on the switch:

```
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:

```
config 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 domain named eaps_1.

```
config eaps eaps_1 mode master
```

The following command example identifies this switch as a transit node for the domain named eaps_1.

```
config 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:

```
config eaps <name> [hellotime <seconds> | failtime <seconds>]
```



NOTE

This command applies 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 its associated `seconds` parameter to specify the amount of time the master node waits before declaring a failed state and opens the logically blocked VLANs on the secondary port. `seconds` must be greater than the configured value for `hellotime`. The default value is three seconds.



NOTE

Increasing the `failtime` value provides more protection against frequent “flapping” between the complete state and the failed state by waiting long enough to receive a health-check packet when the network is congested.

**NOTE**

When the master node declares a failed state, it also flushes its forwarding database (FDB) and sends a “flush FDB” message to all the transit switches on the ring by way of the control VLAN. The reason for flushing the FDB is so that the switches can relearn the new directions to reach Layer 2 end stations via the reconfigured topology.

The following command examples configure the hellotime value for the EAPS domain “eaps_1” to 2 seconds and the failtime value to 10 seconds.

```
config eaps eaps_1 hellotime 2
config eaps eaps_1 failtime 10
```

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:

```
config eaps <name> [primary | secondary] port <port number>
```

The following command example adds port 2 of the switch to the EAPS domain “eaps_1” as the primary port.

```
config eaps eaps_1 primary port 2
```

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**

A control VLAN cannot belong to more than one EAPS domain.

To configure the EAPS control VLAN for the domain, use the following command:

```
config eaps <name> add control vlan <name>
```

**NOTE**

To avoid loops in the network, the control VLAN must NOT be configured with an IP address, and ONLY ring ports may be added to the VLAN.

**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, 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.

**NOTE**

Since the QoS profiles Qp7 and Qp8 share the same hardware queue in the Summit24e3, you must limit the amount of traffic that uses these profiles. Otherwise, the Summit24e3 may drop EAPS control packets, preventing EAPS from operating reliably.

The following command example adds the control VLAN “keys” to the EAPS domain “eaps_1.”

```
config 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:

```
config eaps <name> add protect 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.”

```
config eaps eaps_1 add protect vlan orchid
```

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:

```
unconfig eaps <name> [primary | secondary] port
```

The following command example unconfigures this node's EAPS primary ring port on the domain `eaps_1`:

```
unconfig eaps eaps_1 primary port
```

Displaying EAPS Status Information

To display 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 30 describes the fields and values in the display.

```
* Summit24e3:39 # show eaps detail
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: 13      Port status: Up      Tag status: Tagged
Secondary port: 14    Port status: Up      Tag status: Tagged
Hello Timer interval: 1 sec    Fail Timer interval: 3 sec
```

```

Preforwarding Timer interval: 3 sec
Last update: From Master Id 00:E0:2B:81:20:00, Sat Mar 17 17:03:37 2001
Eaps Domain has following Controller Vlan:
  Vlan Name      VID
  "rhsc"         0020
EAPS Domain has following Protected Vlan(s):
  Vlan Name      VID
  "traffic"      1001
Number of Protected Vlans: 1

```

The following example of the `show eaps {<name>} detail` command displays detailed EAPS information for a single EAPS domain named “eaps2” on the master node. Table 30 describes significant fields and values in the display.

```

* Baker15:4 # show eaps2 detail
Name: "eaps2" (instance=0)
State: Complete      [Running: Yes]
Enabled: Yes      Mode: Master
Primary port: 14      Port status: Up      Tag status: Tagged
Secondary port: 13    Port status: Blocked  Tag status: Tagged
Hello Timer interval: 1 sec    Fail Timer interval: 3 sec
Eaps Domain has following Controller Vlan:
  Vlan Name      VID
  "rhsc"         0020
EAPS Domain has following Protected Vlan(s):
  Vlan Name      VID
  "blue"         1003
  "traffic"      1001
Number of Protected Vlans: 2

```

Table 30: 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. There can only be one EAPS domain on this platform.
EAPSD-Bridge links:	The total number of EAPS bridge links in the system. The maximum count is 255. 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 30: 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. • Complete—The ring is in the COMPLETE state for this EAPS domain. • Failed—There is a break in the ring for this EAPS domain.
[Running: ...]	• Yes—This EAPS domain is running. • No—This EAPS domain is not running.
Enabled:	Indicates whether EAPS is enabled on this domain. • Yes—EAPS is enabled on this domain. • no—EAPS is not enabled.
Mode:	The configured EAPS mode for this switch: transit or master.
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.
Fail Timer interval:	The configured value of the timer.

Table 30: show eaps Display Fields (continued)

Field	Description
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 Vlan:	Lists the assigned name and ID of the control VLAN.
EAPS Domain has ... Protected Vlan: ²	Lists the assigned names and VLAN IDs of all the protected VLANs configured on this EAPS domain.
Number of Protected Vlan:	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.

12

Quality of Service (QoS)

This chapter covers the following topics:

- Overview of Policy-Based Quality of Service on page 133
- Applications and Types of QoS on page 134
- Configuring QoS for a Port or VLAN on page 135
- Traffic Groupings on page 136
 - MAC-Based Traffic Groupings on page 137
 - Explicit Class of Service (802.1p and DiffServ) Traffic Groupings on page 138
 - Physical and Logical Groupings on page 142
- Verifying Configuration and Performance on page 143
- Modifying a QoS Configuration on page 144
- Traffic Rate-Limiting on page 144
- Dynamic Link Context System on page 144

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 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. Each hardware queue can be programmed by ExtremeWare with bandwidth limitation and prioritization parameters. The bandwidth limitation and 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.

Summit24e3 switches support up to four physical queues per port.

**NOTE**

As with all Extreme switch products, 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 31. 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. In addition, RED can be used to reduce session loss if the queue that floods Web traffic becomes over-subscribed.

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 31 summarizes QoS guidelines for the different types of network traffic.

Table 31: 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, RED
File server	Minimum bandwidth

Configuring QoS for a Port or VLAN

Table 32 lists the commands used to configure QoS.

Table 32: QoS Configuration Commands

Command	Description
config ports <portlist> qosprofile <qosprofile>	Configures one or more ports to use a particular QoS profile.
config vlan <name> qosprofile <qosprofile>	Allows you to configure a VLAN to use a particular QoS profile.

Traffic Groupings

Once a QoS profile is modified for bandwidth and priority, you assign traffic a grouping to the profile. 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 -134.

Traffic groupings are separated into the following categories for discussion:

- Access list based information, such as IP source / destination, TCP / UDP port information, and VLANid
- 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 33. 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 33: Traffic Groupings by Precedence

IP Information (Access Lists) Grouping

- Access list precedence determined by user configuration

Explicit Packet Class of Service Groupings

- DiffServ (IP TOS)
- 802.1P

Destination Address MAC-Based Groupings

- Permanent
- Dynamic
- Blackhole

Physical/Logical Groupings

- Source port
 - VLAN
-

Access List Based Traffic Groupings

Access list based traffic groupings are based on any combination of the following items:

- IP source or destination address
- TCP / UDP or other layer 4 protocol
- TCP / UDP port information
- MAC source or destination address
- VLANid

Access list based traffic groupings are defined using access lists. Access lists are discussed in detail in Chapter 9. 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.

MAC-Based Traffic Groupings

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

```
create fdbentry <mac_address> vlan <name> [blackhole | port <portlist> | dynamic]
qosprofile <qosprofile>
```

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

- Permanent
- Dynamic
- Blackhole

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. For example:

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

Dynamic MAC Addresses

Dynamic MAC addresses can be assigned a QoS profile whenever traffic is destined to the MAC address. 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 qosprofile qp3
```

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 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>
```

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. The Summit24e3 switch has the capability of observing and manipulating packet marking information with no performance penalty.

The documented capabilities for 802.1p priority markings or DiffServ capabilities 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.



NOTE


Re-marking DiffServ code points is supported through access lists. See Chapter 9, “Access Policies”, for more information.

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 24.

Figure 24: 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. The Summit24e3 switch

supports four hardware queues. The transmitting hardware queue determines the bandwidth management and 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 34.

Table 34: 802.1p Priority Value-to-QoS Profile to Hardware Queue Default Mapping

Priority Value	QoS Profile	Hardware Queue Priority Value
0	Qp1	1
1	Qp2	1
2	Qp3	2
3	Qp4	2
4	Qp5	3
5	Qp6	3
6	Qp7	4
7	Qp8	4

802.1p Commands

Table 35 shows the command used to configure 802.1p priority. This is explained in more detail in the following paragraphs.

Table 35: 802.1p Configuration Commands

Command	Description
<code>config vlan <name> priority <number></code>	Configures the 802.1p priority value for 802.1Q VLAN tags. The value for <code>priority</code> is an integer between 0 and 7.

Configuring 802.1p Priority

When a packet is transmitted by the switch, you can configure the 802.1p priority field that is placed in the 802.1Q tag. You can configure the priority to be a number between 0 and 7, using the following command:

```
config vlan <name> priority <number>
```

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 Chapter 9, “Access Policies”, 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 in the Summit24e3 switch.

Figure 25 shows the encapsulation of an IP packet header.

Figure 25: IP packet header encapsulation



EW_023

Table 36 lists the commands used to configure DiffServ. Some of the commands are described in more detail in the following paragraphs.

Table 36: DiffServ Configuration Commands

Command	Description
disable diffserv examination ports [<portlist> all]	Disables the examination of the diffserv field in an IP packet.
enable diffserv examination ports [<portlist> all]	Enables the diffserv field of an ingress IP packet to be examined by the switch in order to select a QoS profile. The default setting is disabled.

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 view DiffServ information, use the following command:

```
enable diffserv examination ports [<portlist> | all]
```

**NOTE**

DiffServ examination requires one access mask while it is enabled. See “Maximum Entries” on page 95 for more information.

Changing DiffServ Code point assignments in the QoS Profile

The DiffServ code point has 64 possible values ($2^6 = 64$). By default, the values are grouped and assigned to the default QoS profiles listed in Table 37.

Table 37: Default Code Point-to-QoS Profile Mapping

Code Point	QoS Profile
0-7	Qp1
8-15	Qp2
16-23	Qp3
24-31	Qp4
32-39	Qp5
40-47	Qp6
48-55	Qp7
56-63	Qp8

You can change the QoS profile assignment for a code point by using an access list. See Chapter 9, “Access Policies”, for more information.

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 Chapter 9, “Access Policies”, for more information on using access lists. You will 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 <portlist> info {detail}
```

**NOTE**

The show ports command displays only the default code point mapping.

DiffServ Examples

For information on the access list and access mask commands in the following examples, see Chapter 9, “Access Policies”.

Use the following command to use the DiffServe code point value to assign traffic to the hardware queues:

```
enable diffserv examination ports all
```

In the following example, all the traffic from network 10.1.2.x is assigned the DiffServe code point 23 and the 802.1p value of 2:

```
create access-mask SrcIpMask source-ip/24
create access-list TenOneTwo access-mask SrcIpMask source-ip 10.1.2.0/24 permit qp3
    set code-point 23 set dot1p 2
```

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:

```
config ports <portlist> qosprofile <qosprofile>
```

In the following modular switch example, all traffic sourced from port 7 uses the QoS profile named *qp3* when being transmitted.

```
config 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:

```
config 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:

```
config vlan servnet qosprofile qp4
```

Verifying Physical and Logical Groupings

To verify settings on ports or VLANs, use the following command:

```
show qosprofile <qosprofile>
```

The same information is also available for ports or VLANs using one of the following commands:

```
show ports <portlist> info {detail}
```

or

```
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 incoming packets on a port or ports. The QoS monitor keeps track of the number of frames and the frames per second, sorted by 802.1p value, on each monitored port.

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.

To view real-time switch per-port performance, use the following command:

```
show ports {<portlist>} qosmonitor
```

The QoS monitor rate screen (packets per second), does not display any results for at least five seconds. Once the rate is displayed, it is updated each second.



NOTE

The QoS monitor can display up to four ports at a time.



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.



NOTE

The QoS monitor requires one access mask until it exits. See “Maximum Entries” on page 95 for more information.

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>
```

Displayed information includes:

- QoS profile name
- Minimum bandwidth
- Maximum bandwidth

- Priority
- A list of all traffic groups 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 information including PACE enable/disable information.
- `show vlan` — Displays the QoS profile assignments to the VLAN.
- `show ports <portlist> info {detail}` — Displays information including QoS information for the port.

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 Summit24e3 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 Chapter 9, “Access Policies”.

Dynamic Link Context System

The Dynamic Link Context System (DLCS) is a feature that snoops WINS NetBIOS packets and creates a mapping between a user name, the IP address or MAC address, and the switch/port. Based on the information in the packet, DLCS can detect when an end station boots up or a user logs in or out, and dynamically maps the end station name to the current IP address and switch/port. This information is available for use by ExtremeWare Enterprise Manager (EEM) version 2.1 or later or ExtremeWare EPICenter in setting policies that can be applied to users and can dynamically follow a user's location. DLCS provides you with valuable information on a user's location and associated network attributes. For DLCS to operate within ExtremeWare, the user or end station must allow for automatic DLCS updates.

This feature should only be used in conjunction with the EEM Policy System or ExtremeWare EPICenter Policy System. Refer to the ExtremeWare Enterprise Manager or ExtremeWare EPICenter documentation for more information.

DLCS Guidelines

Follow these guidelines when using DLCS:

- Only one user is allowed on one workstation at a given time.
- A user can be logged into many workstations simultaneously.
- An IP-address can be learned on only one port in the network at a given time.
- Multiple IP-addresses can be learned on the same port.
- DLCS mapping is flushed when a user logs in or logs out, or when an end-station is shutdown.

DLCS Limitations

Consider the following limitations concerning data received from WINS snooping:

- DLCS does not work for the WINS server. This is because the WINS server does not send NETBIOS packets on the network (these packets are address to itself).
- When the IP address of a host is changed, and the host is not immediately rebooted, the old host-to-IP address mapping is never deleted. You must delete the mapping of the host-to-IP address through the EEM Policy Manager or ExtremeWare EPICenter Policy Manager.
- When the host is moved from one port to another port on a switch, the old entry does not age out unless the host is rebooted or a user login operation is performed after the host is moved.
- DLCS information is dynamic, therefore, if the switch is rebooted, the information is lost. This information is still stored in the policy-server. To delete the information from the policy system, you must explicitly delete configuration parameters from the EEM or ExtremeWare EPICenter Policy Applet user interface. As a workaround, you can delete the switch that was rebooted from the list of managed devices in the EEM or EPICenter Inventory Applet, and re-add the switch to the Inventory Manager.
- DLCS is not supported on hosts that have multiple NIC cards.
- IPQoS is not supported to a WINS server that is serving more than one VLAN. If you attempt to add a WINS server to serve more than one VLAN, and there are IPQoS rules defined for that server, the command to add the WINS server is rejected.

DLCS Commands

The DLCS commands are described in Table 38.

Table 38: DLCS Configuration Commands

Command	Description
clear dlcs	Clears learned DLCS data.
disable dlcs	Disables snooping of DLCS packets.
disable dlcs ports <port-number>	Disables port on which DLCS packets are snooped.
enable dlcs	Enables snooping of DLCS packets.
enable dlcs ports <port-number>	Enables port on which DLCS packets are snooped.
show dlcs	Displays ports which are snooping WINS packets, along with the data that has been learned.

13

Status Monitoring and Statistics

This chapter describes the following topics:

- Status Monitoring on page 147
- Port Statistics on page 149
- Port Errors on page 149
- Port Monitoring Display Keys on page 150
- Setting the System Recovery Level on page 151
- Logging on page 151
- RMON on page 155

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.

Status Monitoring

The status monitoring facility provides information about the switch. This information may be useful for your technical support representative if you have a problem. ExtremeWare includes many show commands that display information about different switch functions and facilities.



NOTE

For more information about show commands for a specific ExtremeWare feature, see the appropriate chapter in this guide.

Table 39 describes commands that are used to monitor the status of the switch.

Table 39: Status Monitoring Commands

Command	Description
show diag	Displays software diagnostics.

Table 39: Status Monitoring Commands (continued)

Command	Description
show log {<priority>}	Displays the current snapshot of the log. Options include: • <code>priority</code> — Filters the log to display message with the selected priority or higher (more critical). Priorities include critical, emergency, alert, error, warning, notice, info, and debug. If not specified, all messages are displayed.
show log config	Displays the log configuration, including the syslog host IP address, the priority level of messages being logged locally, and the priority level of messages being sent to the syslog host.
show memory {detail}	Displays the current system memory information. Specify the <code>detail</code> option to view task-specific memory usage.
show switch	Displays the current switch information, including: • <code>sysName</code>, <code>sysLocation</code>, <code>sysContact</code> • MAC address • Current time and time, system uptime, and time zone • Operating environment (fans) • NVRAM configuration information • Scheduled reboot information
show tech-support	Displays the output for the following commands: • <code>show version</code> • <code>show switch</code> • <code>show config</code> • <code>show diag</code> • <code>show gdb</code> • <code>show iparp</code> • <code>show ipfdb</code> • <code>show ipstats</code> • <code>show iproute</code> • <code>show ipmc cache detail</code> • <code>show igmp snooping detail</code> • <code>show memory detail</code> • <code>show log</code> It also displays the output from internal debug commands. This command disables the CLI paging feature.
show version	Displays the hardware and software versions currently running on the switch.

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 <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 <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 <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.
- **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 40 describes the keys used to control the displays that appear when you issue any of the `show port` commands.

Table 40: 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 <code>show port utilization</code> command only.

Setting the System Recovery Level

You can configure the system to automatically reboot after a software task exception, using the following command:

```
config sys-recovery-level [none | critical | all]
```

Where the following is true:

- **none** — Configures the level to recovery without a system reboot.
- **critical** — Configures ExtremeWare to log an error into the syslog and automatically reboot the system after a critical exception.
- **all** — Configures ExtremeWare to log an error into the syslog and automatically reboot the system after any exception.

The default setting is **none**.



NOTE

*Extreme Networks recommends that you set the system recovery level to **critical**. This allows ExtremeWare to log an error to the syslog and automatically reboot the system after a critical exception.*

Logging

The switch 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** — Table 41 describes the three levels of importance that the system can assign to a fault.

Table 41: Fault Levels Assigned by the Switch

Level	Description
Critical	A desired switch function is inoperable. The switch may need to be reset.
Warning	A noncritical error that may lead to a function failure.
Informational	Actions and events that are consistent with expected behavior.
Debug	Information that is useful when performing detailed troubleshooting procedures.

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. To remove log entries of all levels (including warning or critical), use the following command:

```
clear log static
```

- **Subsystem** — The subsystem refers to the specific functional area to which the error refers. Table 42 describes the subsystems.

Table 42: Fault Log Subsystems

Subsystem	Description
Syst	General system-related information. Examples include memory, power supply, security violations, fan failure, overheat condition, and configuration mode.
STP	STP information. Examples include an STP state change.
Brdg	Bridge-related functionality. Examples include low table space and queue overflow.
SNMP	SNMP information. Examples include community string violations.
Telnet	Information related to Telnet login and configuration performed by way of a Telnet session.
VLAN	VLAN-related configuration information.
Port	Port management-related configuration. Examples include port statistics and errors.

- **Message** — The message contains the log information with text that is specific to the problem.

Local Logging

The switch maintains 1,000 messages in its internal log. You can display a snapshot of the log at any time by using the following command:

```
show log {<priority>}
```

where the following is true:

- **priority** — Filters the log to display message with the selected priority or higher (more critical). Priorities include (in order) critical, emergency, alert, error, warning, notice, info, and debug. If not specified, all messages are displayed.

Real-Time Display

In addition to viewing a snapshot of the log, you can configure the system to maintain a running real-time display of log messages on the console. To turn on the log display, use the following command:

```
enable log display
```

To configure the log display, use the following command:

```
config log display {<priority>}
```

If **priority** is not specified, only messages of critical priority are displayed.

If you enable the log display on a terminal connected to the console port, your settings will remain in effect even after your console session is ended (unless you explicitly disable the log display).

When using a Telnet connection, if your Telnet session is disconnected (because of the inactivity timer, or for other reasons), the log display is automatically halted. You must restart the log display by using the `enable log display` command.

Remote Logging

In addition to maintaining an internal log, the switch supports remote logging by way of the UNIX syslog host facility. To enable remote logging, follow these steps:

- 1 Configure the syslog host to accept and log messages.
- 2 Enable remote logging by using the following command:

```
enable syslog
```

- 3 Configure remote logging by using the following command:

```
config syslog {add} <ipaddress> <facility> {<priority>}
```

Specify the following:

- `ipaddress` — The IP address of the syslog host.
- `facility` — The syslog facility level for local use. Options include `local0` through `local7`.
- `priority` — Filters the log to display message with the selected priority or higher (more critical). Priorities include (in order) critical, emergency, alert, error, warning, notice, info, and debug. If not specified, only critical priority messages are sent to the syslog host.



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 are logged to the system log. 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.

Logging Commands

The commands described in Table 43 allow you to configure logging options, reset logging options, display the log, and clear the log.

Table 43: Logging Commands

Command	Description
<code>clear counters</code>	Clears all switch statistics and port counters.
<code>clear log {static}</code>	Clears the log. If <code>static</code> is specified, the critical log messages are also cleared.

Table 43: Logging Commands (continued)

Command	Description
config log display {<priority>}	Configures the real-time log display. Options include: <code>priority</code> — Filters the log to display messages with the selected priority or higher (more critical). Priorities include critical, emergency, error, alert, warning, notice, info, and debug. If not specified, informational priority messages and higher are displayed.
config syslog {add} <host name/ip> <facility> {<priority>}	Configures the syslog host address and filters messages sent to the syslog host. Up to 4 syslog servers can be configured. Options include: <code>host name/ip</code>— The IP address or name of the syslog host. <code>facility</code> — The syslog facility level for local use (local0 - local7). <code>priority</code> — Filters the log to display messages with the selected priority or higher (more critical). Priorities include critical, emergency, alert, error, warning, notice, info, and debug. If not specified, only critical priority messages and are sent to the syslog host.
config syslog delete <host name/ip> <facility> {<priority>}	Deletes a syslog host address. <code>facility</code> — The syslog facility level for local use (local0 - local7). <code>priority</code> — Filters the log to display messages with the selected priority or higher (more critical). Priorities include critical, emergency, alert, error, warning, notice, info, and debug. If not specified, only critical priority messages and are sent to the syslog host.
disable cli-config-logging	Disables configuration logging.
disable log display	Disables the log display.
disable syslog	Disables logging to a remote syslog host.
enable cli-config-logging	Enables the logging of CLI configuration commands to the Syslog for auditing purposes. The default setting is enabled.
enable log display	Enables the log display.
enable syslog	Enables logging to a remote syslog host.
show log {<priority>}	Displays the current snapshot of the log. Options include: <code>priority</code> — Filters the log to display message with the selected priority or higher (more critical). Priorities include critical, emergency, alert, error, warning, notice, info, and debug. If not specified, all messages are displayed.

Table 43: Logging Commands (continued)

Command	Description
show log config	Displays the log configuration, including the syslog host IP address, the priority level of messages being logged locally, and the priority level of messages being sent to the syslog host.

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 the following command:

```
[enable | disable] rmon
```

By default, RMON is disabled. However, even in the disabled state, the switch response 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 44.

Table 44: 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 5, “Managing the Switch”.

14

Spanning Tree Protocol (STP)

This chapter describes the following topics:

- Overview of the Spanning Tree Protocol on page 159
- Spanning Tree Domains on page 159
- STP Configurations on page 160
- Configuring STP on the Switch on page 162
- Displaying STP Settings on page 165
- Disabling and Resetting STP on page 165

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.



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.

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. Once the STPD is created, one or more VLANs can be assigned to it.

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 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



NOTE

Ensure that multiple STPD instances within a single switch do not see each other in the same broadcast domain. This could happen if, for example, another external bridge is used to connect VLANs belonging to separate STPDs.

If you delete an 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.

Defaults

The default device configuration contains a single STPD called *s0*. The default VLAN is a member of STPD *s0*.

All STP parameters default to the IEEE 802.1D values, as appropriate.

STPD BPDU Tunneling

You can configure ExtremeWare to allow a BPDU 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 the following command:

```
[enable | disable] ignore-bpdu 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.

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.

Figure 26 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 the default STPD, but not assigned to either STPD1 or STPD2.

Figure 26: Multiple Spanning Tree Domains




SH_008

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 26, the connection between switch A and switch B 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 not been assigned to either STPD1 or 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 27 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 27: Tag-based STP configuration

SH_009

The tag-based network in Figure 27 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.

Configuring STP on the Switch

To configure STP, follow these steps:

- 1 Create one or more STP domains using the following command:

```
create stpd <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:

```
config stpd <stpd_name> add vlan <name>
```

3 Enable STP for one or more STP domains using the following command:

```
enable stpd {<stpd_name>}
```



NOTE

All VLANs belong to a STPD. If you do not want to run STP on a VLAN, you must add the VLAN to a STPD that is disabled.

Once you have created the STPD, you can optionally configure STP parameters for the STPD.



CAUTION

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

The following parameters can be configured on each port:

- Path cost
- Port priority



NOTE

The device supports the RFC 1493 Bridge MIB. Parameters of only the s0 default STPD are accessible through this MIB.

Table 45 shows the commands used to configure STP.

Table 45: STP Configuration Commands

Command	Description
config stpd <stpd_name> add vlan <name>	Adds a VLAN to the STPD.
config stpd <stpd_name> forwarddelay <value>	Specifies the time (in seconds) that the ports in this STPD spend in the listening and learning states when the switch is the Root Bridge. The range is 4 through 30. The default setting is 15 seconds.
config stpd <stpd_name> hellotime <value>	Specifies the time delay (in seconds) between the transmission of BPDUs from this STPD when it is the Root Bridge. The range is 1 through 10. The default setting is 2 seconds.

Table 45: STP Configuration Commands (continued)

Command	Description
config stpd <stpd_name> maxage <value>	Specifies the maximum age of a BPDU in this STPD. The range is 6 through 40. The default setting is 20 seconds. Note that the time must be greater than, or equal to $2 * (\text{Hello Time} + 1)$ and less than, or equal to $2 * (\text{Forward Delay} - 1)$.
config stpd <stpd_name> ports cost <value> <portlist>	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: - n For a 10 Mbps port, the default cost is 100. - n For a 100 Mbps port, the default cost is 19.
config stpd <stpd_name> ports priority <value> <portlist>	Specifies the priority of the port in this STPD. By changing the priority of the port, you can make it more or less likely to become the root port. The range is 0 through 31. The default setting is 16. A setting of 0 indicates the lowest priority.
config stpd <stpd_name> priority <value>	Specifies the priority of the STPD. By changing the priority of the STPD, you can make it more or less likely to become the root bridge. The range is 0 through 65,535. The default setting is 32,768. A setting of 0 indicates the highest priority.
create stpd <stpd_name>	Creates an STPD. When created, an STPD has the following default parameters: - n Bridge priority — 32,768 - n Hello time — 2 seconds - n Forward delay — 15 seconds
enable ignore-bpdu vlan <name>	Configures the switch to ignore STP BPDUs, which prevents ports in the VLAN from becoming part of an STPD. This command is useful when you have a known topology with switches outside your network, and wish to keep the root bridge within your network. The default setting is disabled.
enable ignore-stp vlan <vlan name>	Configures the switch to ignore the STP protocol, and not block traffic for the VLAN(s). This command is useful when multiple VLANs share the same physical ports, but only some of the VLANs require STP protection. The default setting is disabled.
enable stpd {<stpd_name>}	Enables the STP protocol for one or all STPDs. The default setting is disabled.
enable stpd ports {<portlist>}	Enables the STP protocol on one or more ports. If STPD is enabled for a port, bridge protocol data units (BPDUs) will be generated on that port if STP is enabled for the associated STPD. The default setting is enabled.

STP Configuration Example

The following Summit24e3 switch 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
config stpd backbone_st add vlan manufacturing
enable stpd backbone_st
disable stpd backbone_st port 1-7,12
```

Displaying STP Settings

To display STP settings, use the following command:

```
show stpd {<stpd_name>}
```

This command displays the following information:

- STPD name
- Bridge ID
- STPD configuration information

To display the STP state of a port, use the following command:

```
show stpd <stpd_name> port <portlist>
```

This command displays the following information:

- STPD port configuration
- STPD state (root bridge, and so on)
- STPD port state (forwarding, blocking, and so on)

Disabling and Resetting STP

To disable STP or return STP settings to their defaults, use the commands listed in Table 46.

Table 46: STP Disable and Reset Commands

Command	Description
delete stpd <stpd_name>	Removes an STPD. An STPD can only be removed if all VLANs have been deleted from it. The default STPD, s0, cannot be deleted.
disable ignore-bpdu vlan <name>	Allows the switch to recognize STP BPDUs.
disable ignore-stp vlan <name>	Allows a VLAN to use STP port information.
disable stpd [<stpd_name> all]	Disables the STP mechanism on a particular STPD, or for all STPDs.
disable stpd ports <portlist>	Disables STP on one or more ports. Disabling STP on one or more ports puts those ports in <i>forwarding</i> state; all BPDUs received on those ports will be disregarded.

Table 46: STP Disable and Reset Commands

Command	Description
unconfig stpd {<stpd_name>}	Restores default STP values to a particular STPD or to all STPDs.



IP Unicast Routing

This chapter describes the following topics:

- Overview of IP Unicast Routing on page 167
- Proxy ARP on page 170
- Relative Route Priorities on page 171
- Configuring IP Unicast Routing on page 172
- IP Commands on page 172
- Routing Configuration Example on page 177
- Displaying Router Settings on page 178
- Resetting and Disabling Router Settings on page 178
- Configuring DHCP/BOOTP Relay on page 179
- UDP-Forwarding on page 180

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, refer to Chapter 16.

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 same IP address and subnet on different VLANs.

In Figure 28, a Summit24e3 switch is depicted with two VLANs defined; *Finance* and *Personnel*. Ports 2 and 4 are assigned to *Finance*; ports 3 and 5 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 28: 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 | disable] rip export static
[enable | disable] ospf export 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 (refer to Table 50, later in this chapter)
- Static routes
- Directly attached network interfaces that are not active.



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 eight 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 the following command:

```
[enable | disable] ipforwarding fast-direct-broadcast [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>
```

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:

```
config iparp add proxy <ipaddress> {<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 47 lists the relative priorities assigned to routes depending upon the learned source of the route.



CAUTION

Although these priorities can be changed, do not attempt any manipulation unless you are expertly familiar with the possible consequences.

Table 47: Relative Route Priorities

Route Origin	Priority
Direct	10
BlackHole	50
Static	1100
ICMP	1200
OSPFIntra	2200
OSPFInter	2300
RIP	2400
OSPFExtern1	3200
OSPFExtern2	3300

Table 47: Relative Route Priorities (continued)

Route Origin	Priority
BOOTP	5000

To change the relative route priority, use the following command:

```
config 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:

```
config vlan <name> ipaddress <ipaddress> {<mask>}
```

Ensure that each VLAN has a unique IP address.

- 3 Configure a default route using the following command:

```
config 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 {vlan <name>}
```

- 5 Turn on RIP or OSPF using one of the following commands:

```
enable rip
```

```
enable ospf
```

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.

IP Commands

Table 48 describes the commands used to configure basic IP settings.

Table 48: Basic IP Commands

Command	Description
<code>clear iparp {<ipaddress> <mask> vlan <vlan>}</code>	Removes dynamic entries in the IP ARP table. Permanent IP ARP entries are not affected.
<code>clear ipfdb {<ipaddress> <netmask> vlan <name>}</code>	Removes the dynamic entries in the IP forwarding database. If no options are specified, all dynamic IP FDB entries are removed.
<code>config bootprelay add <ipaddress></code>	Adds the IP destination address to forward BOOTP packets.
<code>config bootprelay delete [<ipaddress> all]</code>	Removes one or all IP destination addresses for forwarding BOOTP packets.
<code>config iparp add <ipaddress> <mac_address></code>	Adds a permanent entry to the ARP table. Specify the IP address and MAC address of the entry.
<code>config iparp add proxy <ipaddress> {<mask>} {<mac_address>} {always}</code>	Configures proxy ARP entries. When <code>mask</code> is not specified, an address with the mask 255.255.255.255 is assumed. When <code>mac_address</code> is not specified, the MAC address of the switch is used in the ARP Response. When <code>always</code> is specified, the switch answers ARP Requests without filtering requests that belong to the same subnet of the receiving router interface.
<code>config iparp delete <ipaddress></code>	Deletes an entry from the ARP table. Specify the IP address of the entry.
<code>config iparp delete proxy [<ipaddress> {<mask>} all]</code>	Deletes one or all proxy ARP entries.
<code>config iparp timeout <minutes></code>	Configures the IP ARP timeout period. The default setting is 20 minutes. A setting of 0 disables ARP aging. The maximum aging time is 32,767 minutes.
<code>disable bootp vlan [<name> all]</code>	Disables the generation and processing of BOOTP packets.
<code>disable bootprelay</code>	Disables the forwarding of BOOTP requests.
<code>disable ipforwarding {vlan <name>}</code>	Disables routing for one or all VLANs.
<code>disable ipforwarding broadcast {vlan <name>}</code>	Disables routing of broadcasts to other networks.
<code>disable loopback-mode vlan [<name> all]</code>	Disables loopback-mode on an interface.
<code>enable bootp vlan [<name> all]</code>	Enables the generation and processing of BOOTP packets on a VLAN to obtain an IP address for the VLAN from a BOOTP server. The default setting is enabled for all VLANs.
<code>enable bootprelay</code>	Enables the forwarding of BOOTP and Dynamic Host Configuration Protocol (DHCP) requests.

Table 48: Basic IP Commands (continued)

Command	Description
enable ipforwarding {vlan <name>}	Enables IP routing for one or all VLANs. If no argument is provided, enables routing for all VLANs that have been configured with an IP address. The default setting for <code>ipforwarding</code> is disabled.
enable ipforwarding broadcast {vlan <name>}	Enables forwarding IP broadcast traffic for one or all VLANs. If no argument is provided, enables broadcast forwarding for all VLANs. To enable, <code>ipforwarding</code> must be enabled on the VLAN. The default setting is disabled.
enable loopback-mode vlan [<name> all]	Enables a loopback mode on an interface. If loopback is enabled, the router interface remains in the UP state, even if no ports are defined in the VLAN. As a result, the subnet is always advertised as one of the available routes.

Table 49 describes the commands used to configure the IP route table.

Table 49: Route Table Configuration Commands

Command	Description
config iproute add <ipaddress> <mask> <gateway> <metric>	Adds a static address to the routing table. Use a value of 255.255.255.255 for <code>mask</code> to indicate a host entry.
config iproute add blackhole <ipaddress> <mask>	Adds a <code>blackhole</code> address to the routing table. All traffic destined for the configured IP address is dropped, and no Internet Control Message Protocol (ICMP) message is generated.
config iproute add default <gateway> {<metric>}	Adds a default gateway to the routing table. A default gateway must be located on a configured IP interface. If no metric is specified, the default metric of 1 is used. Use the unicast-only or multicast-only options to specify a particular traffic type. If not specified, both unicast and multicast traffic uses the default route.
config iproute delete <ipaddress> <mask> <gateway>	Deletes a static address from the routing table.
config iproute delete blackhole <ipaddress> <mask>	Deletes a <code>blackhole</code> address from the routing table.
config iproute delete default <gateway>	Deletes a default gateway from the routing table.
config iproute priority [rip bootp icmp static ospf-intra ospf-inter ospf-as-external ospf-extern1 ospf-extern2] <priority>	Changes the priority for all routes from a particular route origin.
disable iproute sharing	Disables load sharing for multiple routes.

Table 49: Route Table Configuration Commands (continued)

Command	Description
enable iproute sharing	Enables load sharing if multiple routes to the same destination are available. Only paths with the same lowest cost are shared. The default setting is disabled.
rtlookup [<ipaddress> <hostname>]	Performs a look-up in the route table to determine the best route to reach an IP address.

Table 50 describes the commands used to configure IP options and the ICMP protocol.

Table 50: ICMP Configuration Commands

Command	Description
config irdp [multicast broadcast]	Configures the destination address of the router advertisement messages. The default setting is <code>multicast</code> .
config irdp <mininterval> <maxinterval> <lifetime> <preference>	Configures the router advertisement message timers, using seconds. Specify: <code>mininterval</code> — The minimum amount of time between router advertisements. The default setting is 450 seconds. <code>maxinterval</code> — The maximum time between router advertisements. The default setting is 600 seconds. <code>lifetime</code> — The default setting is 1,800 seconds. <code>preference</code> — The preference level of the router. An ICMP Router Discover Protocol (IRDP) client always uses the router with the highest preference level. Change this setting to encourage or discourage the use of this router. The default setting is 0.
disable icmp parameter-problem {vlan <name>}	Disables the generation of ICMP messages for the parameter problem packet type.
disable ip-option loose-source-route	Disables the loose source route IP option.
disable ip-option record-route	Disables the record route IP option.
disable ip-option record-timestamp	Disables the record timestamp IP option.
disable ip-option strict-source-route	Disables the strict source route IP option.
disable ip-option use-router-alert	Disables the generation of the router alert IP option.
enable icmp address-mask {vlan <name>}	Enables the generation of an ICMP address-mask reply (type 18, code 0) when an ICMP address mask request is received. The default setting is enabled. If a VLAN is not specified, the command applies to all IP interfaces.

Table 50: ICMP Configuration Commands (continued)

Command	Description
enable icmp parameter-problem {vlan <name>}	Enables the generation of an ICMP parameter-problem message (type 12) when the switch cannot properly process the IP header or IP option information. The default setting is enabled. If a VLAN is not specified, the command applies to all IP interfaces.
enable icmp port-unreachables {vlan <name>}	Enables the generation of ICMP port unreachable messages (type 3, code 3) when a TCP or UDP request is made to the switch, and no application is waiting for the request, or access policy denies the request. The default setting is enabled. If a VLAN is not specified, the command applies to all IP interfaces.
enable icmp redirects {vlan <name>}	Enables the generation of an ICMP redirect message (type 5) when a packet must be forwarded out on the ingress port. The default setting is enabled. If a VLAN is not specified, the command applies to all IP interfaces.
enable icmp time-exceeded {vlan <name>}	Enables the generation of an ICMP time exceeded message (type 11) when the TTL field expires during forwarding. IP multicast packets do not trigger ICMP time exceeded messages. The default setting is enabled. If a VLAN is not specified, the command applies to all IP interfaces.
enable icmp timestamp {vlan <name>}	Enables the generation of an ICMP timestamp response (type 14, code 0) when an ICMP timestamp request is received. The default setting is enabled. If a VLAN is not specified, the command applies to all IP interfaces.
enable icmp unreachable {vlan <name>}	Enables the generation of ICMP network unreachable messages (type 3, code 0), and host unreachable messages (type 3, code 1) when a packet cannot be forwarded to the destination because of unreachable route or host. ICMP packet processing on one or all VLANs. The default setting is enabled. If a VLAN is not specified, the command applies to all IP interfaces.
enable icmp userredirects	Enables the modification of route table information when an ICMP redirect message is received. This option applies to the switch when it <i>is not configured for routing</i> . The default setting is disabled.
enable ip-option loose-source-route	Enables the loose source route IP option.
enable ip-option record-route	Enables the record route IP option.
enable ip-option record-timestamp	Enables the record timestamp IP option.
enable ip-option strict-source-route	Enables the strict source route IP option.
enable ip-option use-router-alert	Enables the switch to generate the router alert IP option with routing protocol packets.
enable irdp {vlan <name>}	Enables the generation of ICMP router advertisement messages on one or all VLANs. The default setting is enabled.
unconfig icmp	Resets all ICMP settings to the default values.


Table 50: ICMP Configuration Commands (continued)

Command	Description
unconfig irdp	Resets all router advertisement settings to the default values.

Routing Configuration Example

Figure 29 illustrates a Summit24e3 switch that has two VLANs defined as follows:

- *Finance*
 - Contains ports 2 and 4.
 - IP address 192.207.35.1.
- *Personnel*
 - Contains ports 3 and 5.
 - IP address 192.207.36.1.

Figure 29: Unicast routing configuration example

In this configuration, all IP traffic from stations connected to ports 2 and 4 have access to the router by way of the VLAN *Finance*. Ports 3 and 5 reach the router by way of the VLAN *Personnel*.

The example in Figure 29 is configured as follows:

```
create vlan Finance
create vlan Personnel

config Finance add port 2,4
config Personnel add port 3,5
```

```
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
```

Displaying Router Settings

To display settings for various IP routing components, use the commands listed in Table 51.

Table 51: Router Show Commands

Command	Description
show iparp {<ipaddress vlan <name> permanent}	Displays the IP Address Resolution Protocol (ARP) table. You can filter the display by IP address, VLAN, or permanent entries.
show iparp proxy {<ipaddress> {<mask>}}	Displays the proxy ARP table.
show ipconfig {vlan <name>}	Displays configuration information for one or all VLANs.
show ipconfig {vlan <name>} {detail}	Displays IP configuration settings.
show ipfdb {<ipaddress> <netmask> vlan <name> }	Displays the contents of the IP forwarding database (FDB) table. If no option is specified, all IP FDB entries are displayed.
show iproute {priority vlan <vlan> permanent <ipaddress> <netmask> origin [direct static blackhole rip bootp icmp ospf-intra ospf-inter ospf-as-external ospf-extern1 ospf-extern2]} {sorted}	Displays the contents of the IP routing table or the route origin priority.
show ipstats {vlan <name>}	Displays IP statistics for the CPU of the system.

Resetting and Disabling Router Settings

To return router settings to their defaults and disable routing functions, use the commands listed in Table 52.

Table 52: Router Reset and Disable Commands

Command	Description
clear iparp {<ipaddress> vlan <name>}	Removes dynamic entries in the IP ARP table. Permanent IP ARP entries are not affected.
clear ipfdb {<ipaddress> <netmask> vlan <name>]}	Removes the dynamic entries in the IP forwarding database. If no options are specified, all IP FDB entries are removed.
disable bootp vlan [<name> all]	Disables the generation and processing of BOOTP packets.

Table 52: Router Reset and Disable Commands (continued)

Command	Description
<code>disable bootprelay</code>	Disables the forwarding of BOOTP requests.
<code>disable icmp address-mask {vlan <name>}</code>	Disables the generation of an ICMP address-mask reply messages. If a VLAN is not specified, the command applies to all IP interfaces.
<code>disable icmp parameter-problem {vlan <name>}</code>	Disables the generation of ICMP parameter-problem messages. If a VLAN is not specified, the command applies to all IP interfaces.
<code>disable icmp port-unreachables {vlan <name>}</code>	Disables the generation of ICMP port unreachable messages. If a VLAN is not specified, the command applies to all IP interfaces.
<code>disable icmp redirects {vlan <name>}</code>	Disables the generation of ICMP redirect messages. If a VLAN is not specified, the command applies to all IP interfaces.
<code>disable icmp time-exceeded {vlan <name>}</code>	Disables the generation of ICMP time exceeded messages. If a VLAN is not specified, the command applies to all IP interfaces.
<code>disable icmp timestamp {vlan <name>}</code>	Disables the generation of ICMP timestamp response messages. If a VLAN is not specified, the command applies to all IP interfaces.
<code>disable icmp unreachable {vlan <name>}</code>	Disables the generation of ICMP network unreachable messages and host unreachable messages. If a VLAN is not specified, the command applies to all IP interfaces.
<code>disable icmp userredirects</code>	Disables the changing of routing table information when an ICMP redirect message is received.
<code>disable ipforwarding {vlan <name>}</code>	Disables routing for one or all VLANs.
<code>disable ipforwarding broadcast {vlan <name>}</code>	Disables routing of broadcasts to other networks.
<code>disable irdp {vlan <name>}</code>	Disables the generation of router advertisement messages on one or all VLANs.
<code>unconfig icmp</code>	Resets all ICMP settings to the default values.
<code>unconfig irdp</code>	Resets all router advertisement settings to the default values.

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:

```
config bootprelay add <ipaddress>
```

To delete an entry, use the following command:

```
config bootprelay delete {<ipaddress> | all}
```

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
config backbonedhcp add 67 ipaddress 10.1.1.1
config backbonedhcp add 67 ipaddress 10.1.1.2
config labdhcp add 67 vlan labsvrs
config marketing udp-profile backbonedhcp
config operations udp-profile backbonedhcp
config labuser udp-profile labdhcp
```

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 unreachable, port-unreachables, time-exceeded, parameter-problems, redirects, time-stamp, and address-mask requests.

For ICMP packets that are typically routed, you can apply access lists to restrict forwarding behavior. Access lists are described in Chapter 9.

UDP-Forwarding Commands

Table 53 describes the commands used to configure UDP-forwarding.

Table 53: UDP-Forwarding Commands

Command	Description
config udp-profile <profile_name> add <udp_port> [vlan <name> ipaddress <dest_ipaddress>]	Adds a forwarding entry to the specified UDP-forwarding profile name. All broadcast packets sent to <udp_port> are forwarded to either the destination IP address (unicast or subnet directed broadcast) or to the specified VLAN as an all-ones broadcast.
config udp-profile <profile_name> delete <udp_port> [vlan <name> ipaddress <dest_ipaddress>]	Deletes a forwarding entry from the specified udp-profile name.

Table 53: UDP-Forwarding Commands (continued)

Command	Description
config vlan <name> udp-profile <profile_name>	Assigns a UDP-forwarding profile to the source VLAN. Once the UDP profile is associated with the VLAN, the switch picks up any broadcast UDP packets that matches with the user configured UDP port number, and forwards those packets to the user-defined destination. If the UDP port is the DHCP/BOOTP port number, appropriate DHCP/BOOTP proxy functions are invoked.
create udp-profile <profile_name>	Creates a UDP-forwarding profile. You must use a unique name for the UDP-forwarding profile.
delete udp-profile <profile_name>	Deletes a UDP-forwarding profile.
show udp-profile {<profile_name>}	Displays the profile names, input rules of UDP port, destination IP address, or VLAN and the source VLANs to which the profile is applied.
unconfig udp-profile vlan [<name> all]	Removes the UDP-forwarding profile configuration for one or all VLANs.

This chapter describes the following topics:

- Overview on page 183
- Overview of RIP on page 184
- Overview of OSPF on page 186
- Route Re-Distribution on page 191
- Configuring RIP on page 192
- RIP Configuration Example on page 195
- Displaying RIP Settings on page 196
- Resetting and Disabling RIP on page 196
- Configuring OSPF on page 197
- Displaying OSPF Settings on page 202
- Resetting and Disabling OSPF Settings on page 203

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 2328 — *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 for IP unicast routing.

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**

Both 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.

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.

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.

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.

**NOTE**

A Summit24e3 can support up to two non-passive OSPF interfaces, and cannot be a designated or a backup designated router.

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 54 describes LSA type numbers.

Table 54: 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
10	Area scoping
11	AS scoping

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:

```
config ospf ase-limit <number> {timeout <seconds>}
```

where:

- **<number>** – Specifies the number of external LSAs (excluding the default 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.

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.

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:

```
config ospf vlan <name> area <areaid>
```

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 <areaid>
```

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.

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:

```
config ospf area <area_id> nssa {summary | nosummary} stub-default-cost <cost>
{translate}
```

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 30 illustrates a virtual link.



NOTE


Virtual links can not be configured through a stub or NSSA area.

Figure 30: Virtual link using Area 1 as a transit area



Virtual links are also used to repair a discontinuous backbone area. For example, in Figure 31, if the connection between ABR1 and the backbone fails, the connection using ABR2 provides redundancy so that the discontinuous area can continue to communicate with the backbone using the virtual link.

Figure 31: Virtual link providing redundancy



Point-to-Point Support

You can manually configure the OSPF link type for a VLAN. Table 55 describes the link types.

Table 55: 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.

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

Both RIP and OSPF can be enabled simultaneously on the switch. Route re-distribution allows the switch to exchange routes, including static routes, between the two routing protocols. Figure 32 is an example of route re-distribution between an OSPF autonomous system and a RIP autonomous system.

Figure 32: Route re-distribution



EW_019

Configuring Route Re-Distribution

Exporting routes from OSPF to RIP, and from RIP to OSPF, are discreet configuration functions. 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.

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 [static | rip | direct] [cost <metric> [ase-type-1 | ase-type-2]
{tag <number>}]
```

```
disable ospf export [static | rip | direct]
```

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-learned, 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, cost-type, and tag values can be inserted for all the export routes, or route maps can be used for selective insertion.

Verify the configuration using the command:

```
show ospf
```

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:

```
enable rip export [static | direct | ospf | ospf-intra | ospf-inter | ospf-extern1 |
ospf-extern2] cost <metric> tag <number>
```

```
disable rip export [static | direct | ospf | ospf-intra | ospf-inter | ospf-extern1 |
ospf-extern2]
```

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.

OSPF Timers and Authentication

Configuring OSPF timers and authentication on a per-area basis is a shorthand for applying the timers and authentication to each VLAN in the area at the time of configuration. If you add more VLANs to the area, you must configure the timers and authentication for the new VLANs explicitly.

Configuring RIP

Table 56 describes the commands used to configure RIP.

Table 56: RIP Configuration Commands

Command	Description
config rip add vlan [<name> all]	Configures RIP on an IP interface. When an IP interface is created, per-interface RIP configuration is disabled by default.

Table 56: RIP Configuration Commands (continued)

Command	Description
<code>config rip delete vlan [<name> all]</code>	Disables RIP on an IP interface. When RIP is disabled on the interface, the parameters are not reset to their defaults.
<code>config rip garbagetime {<seconds>}</code>	Configures the RIP garbage time. The timer granularity is 10 seconds. The default setting is 120 seconds.
<code>config rip routetimeout {<seconds>}</code>	Configures the route timeout. The default setting is 180 seconds.
<code>config rip rxmode [none v1only v2only any] {vlan <name>}</code>	Changes the RIP receive mode for one or all VLANs. Specify: <code>none</code> — Drop all received RIP packets. <code>v1only</code> — Accept only RIP v1 format packets. <code>v2only</code> — Accept only RIP v2 format packets. <code>any</code> — Accept both RIP v1 and v2 packets. If no VLAN is specified, the setting is applied to all VLANs. The default setting is <code>any</code>.
<code>config rip txmode [none v1only v1comp v2only] {vlan <name>}</code>	Changes the RIP transmission mode for one or all VLANs. Specify: <code>none</code> — Do not transmit any packets on this interface. <code>v1only</code> — Transmit RIP v1 format packets to the broadcast address. <code>v1comp</code> — Transmit RIP v2 format packets to the broadcast address. <code>v2only</code> — 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 <code>v2only</code>.
<code>config rip updatetime {<seconds>}</code>	Changes the periodic RIP update timer. The default setting is 30 seconds.
<code>config rip vlan [<name> all] cost <number></code>	Configures the cost (metric) of the interface. The default setting is 1.
<code>enable rip</code>	Enables RIP. The default setting is disabled.

Table 56: RIP Configuration Commands (continued)


Command	Description
enable rip aggregation	Enables aggregation of subnet information on interfaces configured to send RIP v2 or RIP v2-compatible traffic. The switch summarizes subnet routes to the nearest class network route. The following rules apply when using RIP aggregation: • Subnet routes are aggregated to the nearest class network route when crossing a class boundary. • Within a class boundary, no routes are aggregated. • If aggregation is enabled, the behavior is the same as in RIP v1. • If aggregation is disabled, subnet routes are never aggregated, even when crossing a class boundary. The default setting is disabled.
enable rip export [static direct ospf ospf-intra ospf-inter ospf-extern1 ospf-extern2] metric <metric> {tag <number>}	Enables RIP to redistribute routes from other routing functions. Specify one of the following: • <code>static</code> — Static routes • <code>direct</code> — Interface routes (only interfaces that have IP forwarding enabled are exported) • <code>ospf</code> — All OSPF routes • <code>ospf-intra</code> — OSPF intra-area routes • <code>ospf-inter</code> — OSPF inter-area routes • <code>ospf-extern1</code> — OSPF AS-external route type 1 • <code>ospf-extern2</code> — OSPF AS-external route type 2 The <code>metric</code> range is 0-15. If set to 0, RIP uses the route metric obtained from the route origin.
enable rip originate-default {always} cost <metric> {tag <number>}	Configures a default route to be advertised by RIP if no other default route is advertised. If <code>always</code> is specified, RIP always advertises the default route to its neighbors. If <code>always</code> is not specified, RIP adds a default route if there is a reachable default route in the route table.
enable rip poisonreverse	Enables the split horizon with poison-reverse algorithm for RIP. The default setting is enabled. If you enable poison reverse and split horizon, poison reverse takes precedence.
enable rip splithorizon	Enables the split horizon algorithm for RIP. Default setting is enabled.
enable rip triggerupdates	Enables triggered updates. <i>Triggered updates</i> are a mechanism for immediately notifying a router's neighbors when the router adds or deletes routes, or changes the metric of a route. The default setting is enabled.

RIP Configuration Example

Figure 33 illustrates a switch that has two VLANs defined as follows:

- *Finance*
 - Contains ports 2 and 4
 - IP address 192.207.35.1.
- *Personnel*
 - Contains ports 3 and 5
 - IP address 192.207.36.1.

Figure 33: RIP configuration example



In this configuration, all IP traffic from stations connected to ports 2 and 4 have access to the router by way of the VLAN *Finance*. Ports 3 and 5 reach the router by way of the VLAN *Personnel*.

The example in Figure 33 is configured as follows:

```
create vlan Finance
create vlan Personnel

config Finance add port 2,4
config Personnel add port 3,5

config Finance ipaddress 192.207.35.1
config Personnel ipaddress 192.207.36.1
```

```
enable ipforwarding
config rip add vlan all
enable rip
```

Displaying RIP Settings

To display settings for RIP, use the commands listed in Table 57.

Table 57: RIP Show Commands

Command	Description
show rip {detail}	Displays RIP configuration and statistics for all VLANs.
show rip stat {detail}	Displays RIP-specific statistics for all VLANs.
show rip stat vlan <name>	Displays RIP-specific statistics for a VLAN.
show rip vlan <name>	Displays RIP configuration and statistics for a VLAN.

Resetting and Disabling RIP

To return RIP settings to their defaults, or to disable RIP, use the commands listed in Table 58.

Table 58: RIP Reset and Disable Commands

Command	Description
config rip delete [vlan <name> all]	Disables RIP on an IP interface. When RIP is disabled on the interface, the parameters are not reset to their defaults.
disable rip	Disables RIP.
disable rip aggregation	Disables the RIP aggregation of subnet information on a RIP v2 interface.
disable rip export [static direct ospf ospf-intra ospf-inter ospf-extern1 ospf-extern2] metric <metric> {tag <number>}	Disables the distribution of non-RIP routes into the RIP domain.
disable rip originate-default	Disables the advertisement of a default route.
disable rip poisonreverse	Disables poison reverse.
disable rip splithorizon	Disables split horizon.
disable rip triggerupdates	Disables triggered updates.
unconfig rip {vlan <name>}	Resets all RIP parameters to match the default VLAN. Does not change the enable/disable state of the RIP settings. If no VLAN is specified, all VLANs are reset.

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.

Table 59 describes the commands used to configure OSPF.

Table 59: OSPF Configuration Commands

Command	Description
<code>config ospf add vlan <name> area <areaid> link-type [auto broadcast point-to-point] {passive}</code>	Configures the OSPF link type. Specify one of the following: <code>auto</code>—ExtremeWare automatically determines the OSPF link type based on the interface type. <code>broadcast</code>—Broadcast link, such as Ethernet. Routers must elect a DR and a BDR during synchronization. <code>point-to-point</code>—Point-to-point link type, such as PPP. The default setting is <code>auto</code>. The <code>passive</code> parameter indicates that the interface does not send or receive OSPF packets.
<code>config ospf vlan <name> neighbor add <ipaddress></code>	Configures the IP address of a point-to-point neighbor.
<code>config ospf vlan <name> neighbor delete <ipaddress></code>	Deletes the IP address of a point-to-point neighbor.
<code>config ospf [area <areaid> vlan [<name> all]] cost [automatic <number>]</code>	Configures the cost metric of one or all VLAN(s). If an area is specified, the cost metric is applied to all VLANs currently within that area. When <code>automatic</code> is specified, the advertised cost is determined from the OSPF metric table and corresponds to the active highest bandwidth port in the VLAN.
<code>config ospf [vlan <name> area <areaid> virtual-link <routerid> <areaid>] authentication [simple-password <password> md5 <md5_key_id> <md5_key>] none encrypted [simple-password <password> md5 <md5_key_id> <md5_key>]</code>	Specifies the authentication password (up to eight characters) or Message Digest 5 (MD5) key for one or all interfaces (VLANs) in an area. The <code>md5_key</code> is a numeric value with the range 0 to 65,536. When the OSPF area is specified, authentication information is applied to all OSPF interfaces within the area.

Table 59: OSPF Configuration Commands (continued)

Command	Description
<pre>config ospf [vlan <name> area <areaid> virtual-link <routerid> <areaid>] timer <retransmission_interval> <transmission_delay> <hello_interval> <dead_interval></pre>	Configures the timers for one interface or all interfaces in the same OSPF area. The following default, minimum, and maximum values (in seconds) are used: - retransmission_interval Default: 5 Minimum: 0 Maximum: 3,600 - transmission_delay Default: 1 Minimum: 0 Maximum: 3,600 - hello_interval Default: 10 Minimum: 1 Maximum: 65,535 - dead_interval Default: 40 Minimum: 1 Maximum: 2,147,483,647
<pre>config ospf add virtual-link <routerid> <areaid></pre>	Adds a virtual link to another ABR. Specify the following: - routerid — Far-end router interface number. - areaid — Transit area used for connecting the two end-points.
<pre>config ospf add vlan <name> area <areaid> {passive}</pre>	Enables OSPF on one or all VLANs (router interfaces). The <areaid> specifies the area to which the VLAN is assigned. The <code>passive</code> parameter indicates that the interface does not send or receive OSPF packets.
<pre>config ospf area <areaid> add range <ipaddress> <mask> [advertise noadvertise] [type 3 type 7]</pre>	Configures a range of IP addresses in an OSPF area. If advertised, the range is exported as a single LSA by the ABR.
<pre>config ospf area <areaid> delete range <ipaddress> <mask></pre>	Deletes a range of IP addresses in an OSPF area.
<pre>config ospf area <areaid> normal</pre>	Configures an OSPF area as a normal area. The default setting is <code>normal</code>.
<pre>config ospf area <areaid> nssa [summary nosummary] stub-default-cost <cost> {translate}</pre>	Configures an OSPF area as a NSSA.
<pre>config ospf area <areaid> stub [summary nosummary] stub-default-cost <cost></pre>	Configures an OSPF area as a stub area.
<pre>config ospf asbr-filter [<access_profile> none]</pre>	Configures a route filter for non-OSPF routes exported into OSPF. If <code>none</code> is specified, no RIP and static routes are filtered.
<pre>config ospf ase-limit <number> {timeout <seconds>}</pre>	Configures OSPF database overflow.
<pre>config ospf ase-summary add <ipaddress> <mask> cost <cost> {<tag_number>}</pre>	Configures an aggregated OSPF external route using the IP addresses specified.

Table 59: OSPF Configuration Commands (continued)

Command	Description
config ospf ase-summary delete <ipaddress> <mask>	Deletes an aggregated OSPF external route.
config ospf delete virtual-link <routerid> <areaid>	Removes a virtual link.
config ospf delete vlan [<name> all]	Disables OSPF on one or all VLANs (router interfaces).
config ospf direct-filter [<access_profile> none]	Configures a route filter for direct routes. If none is specified, all direct routes are exported if <code>ospf export direct</code> is enabled.
config ospf lsa-batching-timer <timer_value>	Configures the OSPF LSA batching timer value. The range is between 0 (disabled) and 600 seconds, using multiples of 5 seconds. The LSAs added to the LSDB during the interval are batched together for refresh or timeout. The default setting is 30 seconds.
config ospf metric-table <10M_cost> <100M_cost> <1G_cost>	Configures the automatic interface costs for 10 Mbps, 100 Mbps, and 1 Gbps interfaces. The default cost for 10 Mbps is 10, for 100 Mbps is 5, and for 4 Gbps is 1.
config ospf routerid [automatic <routerid>]	Configures the OSPF router ID. If automatic is specified, the switch uses the largest IP interface address as the OSPF router ID. The default setting is <code>automatic</code> .
config ospf spf-hold-time {<seconds>}	Configures the minimum number of seconds between Shortest Path First (SPF) recalculations. The default setting is 3 seconds.
config ospf vlan <name> area <areaid>	Changes the area ID of an OSPF interface (VLAN).

Table 59: OSPF Configuration Commands (continued)

Command	Description
<pre>config ospf vlan <vlan> timer <rxmtinterval> <transitdelay> <hellointerval> <routerdeadinterval> [<waitinterval>]</pre>	Configures the OSPF wait interval. Specify the following: <code>rxmtinterval</code> — 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. <code>transitdelay</code> — The length of time it takes to transmit an LSA packet over the interface. The transit delay must be greater than 0. <code>hellointerval</code> — The interval at which routers send hello packets. Smaller times allow routers to discover each other more quickly, but also increase network traffic. The default value is 10 seconds. <code>routerdeadinterval</code> — The interval after which a neighboring router is declared down due to the fact that hello packets are no longer received from the neighbor. This interval should be a multiple of the hello interval. The default value is 40 seconds. <code>waitinterval</code> — The interval between the interface coming up and the election of the DR and BDR. This interval is required by the OSPF standard to be equal to the <code>routerdeadinterval</code>. Under some circumstances, setting the <code>waitinterval</code> to smaller values can help OSPF routers on a broadcast network to synchronize more quickly at the expense of possibly electing an incorrect DR or BDR. This value should not be set to less than the <code>hellointerval</code>. The default value is equal to the <code>routerdeadinterval</code>.
<pre>create ospf area <areaid></pre>	Creates an OSPF area. Area 0 does not need to be created. It exists by default.
<pre>disable ospf capability opaque-lsa</pre>	Disables OSPF opaque LSA support.
<pre>enable ospf</pre>	Enables OSPF process for the router.
<pre>enable ospf capability opaque-lsa</pre>	Enables OSPF opaque LSA support.
<pre>enable ospf export direct [cost <metric> [ase-type-1 ase-type-2] {tag <number>}]</pre>	Enables the distribution of local interface (direct) routes into the OSPF domain. Once enabled, the OSPF router is considered to be an ASBR. The default tag number is 0. The default setting is disabled. Interface routes which correspond to the interface that has OSPF enabled are ignored.
<pre>enable ospf export rip [cost <metric> [ase-type-1 ase-type-2] {tag <number>}]</pre>	Enables the distribution of RIP routes into the OSPF domain. Once enabled, the OSPF router is considered to be an ASBR. The default tag number is 0. The default setting is disabled.

Table 59: OSPF Configuration Commands (continued)

Command	Description
<code>enable ospf export static [cost <metric> [ase-type-1 ase-type-2] {tag <number>}]</code>	Enables the distribution of static routes into the OSPF domain. Once enabled, the OSPF router is considered to be an ASBR. The default tag number is 0. The default setting is disabled.
<code>enable ospf originate-default {always} cost <metric> [ase-type-1 ase-type-2] {tag <number>}</code>	Configures a default external LSA to be generated by OSPF, if no other default route is originated by OSPF by way of RIP and static route re-distribution. If <code>always</code> is specified, OSPF always advertises the default route. If <code>always</code> is not specified, OSPF adds the default LSA if there is a reachable default route in the route table.

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:

```
config ospf vlan <vlan> timer <rxmtinterval> <transitdelay> <hellointerval>
<routerdeadinterval> [<waitinterval>]
```

You can configure the following parameters:

- Retransmit interval (RxmtInterval) — 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.
- Transit delay (TransitDelay) — The length of time it takes to transmit an LSA packet over the interface. The transit delay must be greater than 0.
- Hello interval (HelloInterval) — The interval at which routers send hello packets. Smaller times allow routers to discover each other more quickly, but also increase network traffic. The default value is 10 seconds.
- Dead router wait interval (RouterDeadInterval) — The interval after which a neighboring router is declared down due to the fact that hello packets are no longer received from the neighbor. This interval should be a multiple of the hello interval. The default value is 40 seconds.
- Router wait interval (WaitInterval) — The interval between the interface coming up and the election of the DR and BDR. This interval should be greater than the hello interval. If it is close to the hello interval, the network synchronizes very quickly, but might not elect the correct DR or BDR. The default value is equal to the dead router wait interval.



NOTE

The OSPF standard specifies that wait times are equal to the dead router wait interval.

Displaying OSPF Settings

To display settings for OSPF, use the commands listed in Table 60.

Table 60: OSPF Show Commands

Command	Description
<code>show ospf</code>	Displays global OSPF information.
<code>show ospf area {detail}</code>	Displays information about all OSPF areas.
<code>show ospf area <areaid></code>	Displays information about a particular OSPF area.
<code>show ospf ase-summary</code>	Displays the OSPF external route aggregation configuration.
<code>show ospf interfaces {detail}</code>	Displays information about all OSPF interfaces.
<code>show ospf interfaces {vlan <name> area <areaid>}</code>	Displays information about one or all OSPF interfaces.
<code>show ospf lsdb {detail} area [<areaid> all]</code> <code>[router network summary-net summary-asb as-external external-type7 all]</code>	Displays a table of the current LSDB. You can filter the display using the area ID and LSA type. The default setting is <code>all</code> with no detail. If <code>detail</code> is specified, each entry includes complete LSA information.
<code>show ospf virtual-link {<areaid> <routerid> }</code>	Displays virtual link information about a particular router or all routers.

OSPF LSD 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 [detail | summary | stats] area [all | <areaid>[/<len>]] lstype [all |
as-external | external-type7 | network | router | summary-asb | summary-net] [lsid
<id>[/<len>]] [routerid <id>[/<len>]]
```

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.

Resetting and Disabling OSPF Settings

To return OSPF settings to their defaults, use the commands listed in Table 61.

Table 61: OSPF Reset and Disable Commands

Command	Description
delete ospf area [<areaid> all]	Deletes an OSPF area. Once an OSPF area is removed, the associated OSPF area and OSPF interface information is removed. The backbone area cannot be deleted. A non-empty area cannot be deleted.
disable ospf	Disables OSPF process in the router.
disable ospf export direct	Disables exporting of local interface (direct) routes into the OSPF domain.
disable ospf export rip	Disables exporting of RIP routes in the OSPF domain.
disable ospf export static	Disables exporting of statically configured routes into the OSPF domain.
unconfig ospf {vlan <name> area <areaid>}	Resets one or all OSPF interfaces to the default settings.

17

IP Multicast Routing

This chapter describes the following topics:

- Overview on page 205
- Configuring IP Multicasting Routing on page 206
- Displaying IP Multicast Routing Settings on page 207
- Deleting and Resetting IP Multicast Settings on page 208

For more information on IP multicasting, see the following publications:

- RFC 1112 – *Host Extension for IP Multicasting*
- RFC 2236 – *Internet Group Management Protocol, Version 2*

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 method for the IP host to communicate its multicast group membership to a router (for example, Internet Group Management Protocol (IGMP)).



You should configure IP unicast routing before you configure IP multicast routing.

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 period IGMP query packets. IGMP query 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. The feature reduces the flooding of IP multicast traffic. IGMP snooping optimizes the usage of network bandwidth, and prevents multicast traffic from being flooded to parts of the 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 you are using multicast routing, IGMP snooping must be enabled. If IGMP snooping is disabled, all IGMP and IP multicast traffic floods within a given VLAN. IGMP snooping expects at least one device in the network to periodically generate IGMP query messages. Without an IGMP querier, the switch stops forwarding IP multicast packets to any port.

When a port sends an IGMP leave message, the switch removes the IGMP snooping entry after 10 seconds. 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, the router does not receive any responses to the query, and the router immediately removes the VLAN from the multicast group.

IGMP configuration commands can be found in Table 62.

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 <name>}
```

The `enable ipmcforwarding` command enables IP multicast forwarding on an IP interface. If no options are specified, all configured IP interfaces are affected. By default, `ipmcforwarding` is disabled when you add a new IP interface.

Table 62 describes the commands used to configure the Internet Gateway Message Protocol (IGMP).

Table 62: IGMP Configuration Commands

Command	Description
<pre>config igmp <query_interval> <query_response_interval> <last_member_query_interval></pre>	Configures the IGMP timers. Timers are based on RFC2236. Specify the following: <code>query_interval</code> — The amount of time, in seconds, the system waits between sending out General Queries. The range is 1 to 2,147,483,647 seconds (68 years). The default setting is 125 seconds. <code>query_response_interval</code> — The maximum response time inserted into the periodic General Queries. The range is 1 to 25 seconds. The default setting is 10 seconds. <code>last_member_query_interval</code> — The maximum response time inserted into a Group-Specific Query sent in response to a Leave group message. The range is 1 to 25 seconds. The default setting is 1 second.
<pre>config igmp snooping <router_timeout> <host_timeout></pre>	Configures the IGMP snooping timers. Timers should be set to approximately 2.5 times the router query interval in use on the network. Specify the following: <code>router_timeout</code> — The interval, in seconds, between the last time the router was discovered and the current time. The range is 10 to 2,147,483,647 seconds (68 years). The default setting is 260 seconds. <code>host_timeout</code> — The interval, in seconds, between the last IGMP group report message from the host and the current time. The range is 10 to 2,147,483,647 seconds (68 years). The default setting is 260 seconds.
<pre>enable igmp {vlan <name>}</pre>	Enables IGMP on a router interface. If no VLAN is specified, IGMP is enabled on all router interfaces. The default setting is enabled.
<pre>enable igmp snooping {forward-mcrouter-only}</pre>	Enables IGMP snooping on the switch. If <code>forward-mcrouter-only</code> is specified, the switch forwards all multicast traffic to the multicast router, only. Otherwise, the switch forwards all multicast traffic to any IP router.

Displaying IP Multicast Routing Settings

To display settings for IP multicast routing components, use the commands listed in Table 63.

Table 63: IP Multicast Routing Show Commands

Command	Description
<pre>show igmp snooping {vlan <name>} {detail}</pre>	Displays IGMP snooping registration information, and a summary of all IGMP timers and states.

Table 63: IP Multicast Routing Show Commands (continued)


Command	Description
show ipmc cache {detail} {<group>} {<source> <netmask>}}	Displays the IP multicast forwarding cache.

Deleting and Resetting IP Multicast Settings

To return IP multicast routing settings to their defaults and disable IP multicast routing functions, use the commands listed in Table 64.

Table 64: IP Multicast Routing Reset and Disable Commands

Command	Description
clear igmp snooping {vlan <name>}	Removes one or all IGMP snooping entries.
clear ipmc cache {<group> {<source> <netmask>}}	Resets the IP multicast cache table. If no options are specified, all IP multicast cache entries are flushed.
disable igmp {vlan <name>}	Disables the router-side IGMP processing on a router interface. No IGMP query is generated, but the switch continues to respond to IGMP queries received from other devices. If no VLAN is specified, IGMP is disabled on all router interfaces.
disable igmp snooping	Disables IGMP snooping. IGMP snooping can be disabled only if IP multicast routing is not being used. Disabling IGMP snooping allows all IGMP and IP multicast traffic to flood within a given VLAN.
disable ipmcforwarding {vlan <name>}	Disables IP multicast forwarding.
unconfig igmp	Resets all IGMP settings to their default values and clears the IGMP group table.



Safety Information

Important Safety Information



WARNING!

Read the following safety information thoroughly before installing your Extreme Networks switch. Failure to follow this safety information can lead to personal injury or damage to the equipment.

Installation, maintenance, removal of parts, and removal of the unit and components must be done by qualified service personnel only.

Service personnel are people having appropriate technical training and experience necessary to be aware of the hazards to which they are exposed when performing a task and of measures to minimize the danger to themselves or other people.

Install the unit only in a temperature- and humidity-controlled indoor area free of airborne materials that can conduct electricity. Too much humidity can cause a fire. Too little humidity can produce electrical shock and fire.



NOTE

For more information about the Summit24e3 temperature and humidity ranges, see Appendix B.

Power

The Summit24e3 switch has one power input on the switch.

- The unit must be grounded. Do not connect the power supply unit to an AC outlet without a ground connection.
- The unit must be connected to a grounded outlet to comply with European safety standards.
- The socket outlet must be near the unit and easily accessible. You can only remove power from the unit by disconnecting the power cord from the outlet.
- This unit operates under Safety Extra Low Voltage (SELV) conditions according to IEC 950. The conditions are only maintained if the equipment to which it is connected also operates under SELV conditions.

- The appliance coupler (the connector to the unit and not the wall plug) must have a configuration for mating with an EN60320/IEC320 appliance inlet.
- *France and Peru only*
This unit cannot be powered from IT[†] supplies. If your supplies are of IT type, this unit must be powered by 230 V (2P+T) via an isolation transformer ratio 1:1, with the secondary connection point labeled Neutral, connected directly to ground.

Power Cord

The power cord must be approved for the country where it is used:

- USA and Canada
 - The cord set must be UL-listed and CSA-certified.
 - The minimum specification for the flexible cord is No. 18 AWG (1.5 mm²), Type SVT or SJT, 3-conductor.
 - The cord set must have a rated current capacity of at least the amount rated for each specific product.
 - The AC attachment plug must be an Earth-grounding type with a NEMA 5-15P (10 A, 125 V) configuration.
- Denmark
 - The supply plug must comply with section 107-2-D1, standard DK2-1a or DK2-5a.
- Switzerland
 - The supply plug must comply with SEV / ASE 1011.
- Argentina
 - The supply plug must comply with Argentinian standards.

Connections

Fiber Optic ports - Optical Safety. Never look at the transmit LED/laser through a magnifying device while it is powered on. Never look directly at the fiber port or fiber cable ends when they are powered on.

This is a Class 1 laser device.



WARNING!

Use only for data communications applications that require optical fiber. Use only with the appropriate connector. When not in use, replace dust cover. Using this module in ways other than those described in this manual can result in intense heat that can cause fire, property damage, or personal injury.

Lithium Battery


The battery in the bq4830/DS1644 device is encapsulated and not user-replaceable.

If service personnel disregard the instructions and attempt to replace the bq4830/DS1644, replace the lithium battery with the same or equivalent type, as recommended by the manufacturer.

**WARNING!**

Danger of explosion if battery is incorrectly replaced. Replace only with the same or equivalent type recommended by the manufacturer. Dispose of used batteries according to the manufacturer's instructions.

- Disposal requirements vary by country and by state.
- Lithium batteries are not listed by the Environmental Protection Agency (EPA) as a hazardous waste. Therefore, they can typically be disposed of as normal waste.
- If you are disposing of large quantities, contact a local waste-management service.
- No hazardous compounds are used within the battery module.
- The weight of the lithium contained in each coin cell is approximately 0.035 grams.
- Two types of batteries are used interchangeably:
 - CR chemistry uses manganese dioxide as the cathode material.
 - BR chemistry uses poly-carbonmonofluoride as the cathode material.



Technical Specifications

Physical Dimensions	Summit24e3 Height: 1.75 inches (4.44 cm) Width: 17 inches (43.18 cm) Depth: 8 inches (20.32 cm) Weight: 8 lbs (3.6 kg)
Safety	
Agency Certifications	UL 1950 3rd Edition, listed EN60950:1992/A1-4:1997 plus ZB/ZC Deviations IEC 950CB Low Voltage Directive (LVD) CSA 22.2#950-95 AS/NZS 3260 EN60825-1 FCC CFR 21
Electromagnetic Compatibility	FCC CFR 47 part 15 Class A ICES-0003 A/C108.8-M1983 Class A VCCI Class A AS/NZS 3548 EN55022:1998 Class A CISPR 22:1998 Class A EN55024:1998 includes IEC 61000-4-2, 3, 4, 5, 6, 8, 11 EN 61000-3-2, 3 CNS 13438 Class A
Heat Dissipation	74 W maximum (252 BTU/hr maximum)
Power Supply	
AC Line Frequency	50 Hz to 60 Hz
Input Voltage Options	90 VAC to 264 VAC, auto-ranging
Current Rating	100-120/200-240 VAC 2.0/1.0 A
Switch Power-Off	
Temperature power-off	Summit 24e3 Digital supplies, not Rev. C1: Not drifting: 65° to 70° C (149° to 158° F) Drifting: 50° C (122° F) Digital supplies, Rev. C1: 70° to 75° C (158° to 167° F) Power-One supplies, Rev. OL and earlier: 60° to 65° C (140° to 149° F) Power-One supplies, Rev. OM and later: 75° C (167° F)

Environmental Requirements

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

Certification Marks



CE (European Community)



TUV/GS (German Notified Body)



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 Standards

The following is a list of software standards supported by ExtremeWare for the Summit24e3 switch.

Standards and Protocols	
RFC 1058 RIP	RFC 783 TFTP
RFC 1723 RIP v2	RFC 1542 BootP
RFC 1112 IGMP	RFC 854 Telnet
RFC 2236 IGMP v2	RFC 768 UDP
RFC 2328 OSPF v2 (incl. MD5 authentication)	RFC 791 IP
RFC 2154 OSPF with Digital Signatures (password, MD-5)	RFC 792 ICMP
RFC 1587-NSSA option	RFC 793 TCP
RCC 1765 OSPF Database Overflow	RFC 826 ARP
RFC 2370 OSPF Opaque LSA Option	RFC 2068 HTTP
RFC 1122 Host requirements	RFC 2131 BootP/DHCP relay
IEEE 802.1D-1998 (802.1p) Packet priority	RFC 2030 - Simple Network Time Protocol
IEEE 802.1Q VLAN tagging	RFC 1256 Router discovery protocol
RFC 2474 DiffServ Precedence	RFC 1812 IP router requirement
	RFC 1519 CIDR
Management and Security	
RFC 1157 SNMP v1/v2c	RFC 2239 802.3 MAU MIB
RFC 1213 MIB II	RFC 1724 RIP v2 MIB
RFC 1354 IP forwarding table MIB	RFC 1850 OSPF v2 MIIB
RFC 1493 Bridge MIB	ExtremeWare Enterprise MIB
RFC 2037 Entity MIB	HTML and Telnet management
RFC 1573 Evolution of Interface	RFC 2138 RADIUS
RFC 1643 Ethernet MIB	RFC 2925 Ping MIB
RFC 1757 Four groups of RMON	RFC 2233 Interface MIB
ExtremeWare VLAN Configuration private MIB	RFC 2096 IP Forwarding Table MIB
RFC 2021 RMON probe configuration	999 local messages, criticals stored across reboots



Software Upgrade and Boot Options

This appendix describes the following topics:

- Downloading a New Image on page 217
- Saving Configuration Changes on page 219
- Using TFTP to Upload the Configuration on page 219
- Using TFTP to Download the Configuration on page 221
- Upgrading and Accessing BootROM on page 222
- Boot Option Commands on page 223

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. 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).
- Download the new image to the switch using the following command:

```
download image [<ipaddress> | <hostname>] <filename> {primary | secondary}
```

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 filename of the new image.
`primary` — Indicates the primary image.
`secondary` — Indicates the 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 primary image space is used.

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.

If you have made a mistake, or you must revert to the configuration as it was before you started making changes, you can tell the switch to use the secondary configuration on the next reboot.

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:

```
unconfig 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:

```
unconfig 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 [<ipaddress> | <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

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 config` 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 following command:

```
download configuration [<hostname> | <ipaddress>] <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 [<hostname> | <ipaddress>] <filename> {incremental}
```

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:

```
config download server [primary | secondary] [<hostname> | <ipaddress>] <filename>
```

To enable scheduled incremental downloads, use the following command:

```
download configuration every <hour (0-23)>
```

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 [<hostname> | <ipaddress>] <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

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.

Boot Option Commands


Table 65 lists the CLI commands associated with switch boot options.

Table 65: Boot Option Commands

Command	Description
<code>config download server [primary secondary] [<hostname> <ipaddress>] <filename></code>	Configures the TFTP server(s) used by a scheduled incremental configuration download.
<code>download bootrom [<hostname> <ipaddress>] <filename></code>	Downloads a BOOT ROM image from a TFTP server. The downloaded image replaces the BOOT ROM in the onboard FLASH memory.
 NOTE <i>If this command does not complete successfully, it could prevent the switch from booting.</i>	
<code>download configuration [<hostname> <ipaddress>] <filename> {incremental}</code>	Downloads a complete configuration. Use the <code>incremental</code> keyword to specify an incremental configuration download.
<code>download configuration cancel</code>	Cancels a previously scheduled configuration download.
<code>download configuration every <hour></code>	Schedules a configuration download. Specify the hour using a 24-hour clock, where the range is 0 to 23.

Table 65: Boot Option Commands (continued)

Command	Description
download image [<ipaddress> <hostname>] <filename> {primary secondary}	Downloads a new image from a TFTP server over the network. If no parameters are specified, the image is saved to the current image.
reboot {time <date> <time> cancel}	Reboots the switch at the date and time specified. If you do not specify a reboot time, the reboot happens immediately following the command, and any previously scheduled reboots are cancelled. To cancel a previously scheduled reboot, use the <code>cancel</code> option.
save {configuration} {primary secondary}	Saves the current configuration to nonvolatile storage. You can specify the primary or secondary configuration area. If not specified, the configuration is saved to the primary configuration area.
show configuration	Displays the current configuration to the terminal. You can then capture the output and store it as a file.
upload configuration [<ipaddress> <hostname>] <filename> {every <time>}	Uploads the current run-time configuration to the specified TFTP server. If <code>every <time></code> is specified, the switch automatically saves the configuration to the server once per day, at the specified time. If the time option is not specified, the current configuration is immediately uploaded.
upload configuration cancel	Cancels a previously schedule configuration upload.
use configuration [primary secondary]	Configures the switch to use a particular configuration on the next reboot. Options include the primary configuration area or the secondary configuration area.
use image [primary secondary]	Configures the switch to use a particular image on the next reboot.



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 amber:

The device has failed its Power On Self Test (POST) and you should contact your supplier for advice.

A link is connected, but the Port 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.

Both sides of the Gigabit link must be enabled or disabled. If the two 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 port config
```

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.

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, no flow control.

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 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 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.

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 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 port rx` 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 `config 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 Mini-GBIC. 1000BASE-SX does not work with single-mode fiber (SMF).

VLANs

You cannot add a port to a VLAN:

If you attempt to add a port to the “default” VLAN and get an error message similar to

```
Summit24e3:28 # config vlan default add port 1
ERROR: There is a protocol conflict with adding port 1 untagged to VLAN default
```

you already have a VLAN using untagged traffic on this 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 <name>
```

The solution for this error is to remove port 1 from the VLAN currently using untagged traffic on the port. If this were the “default” VLAN, the command would be

```
Summit24e3:30 # config vlan default del port 1
```

which should now allow you to re-enter the previous command without error as follows:

```
Summit24e3:31 # config vlan red add port 1
```

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.

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

ExtremeWare includes a debug-tracing facility for the switch. The show debug-tracing command can be applied to one or all VLANs, as follows:

```
show debug-tracing {vlan <name>}
```

The debug commands should only be used under the guidance of Extreme Networks technical personnel.

TOP Command

The `top` command is a utility that indicates CPU utilization by process.

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/extreme/support/techsupport.asp>

to download software updates (requires a service contract) and documentation.



Index

Numerics

802.1p configuration commands (table) 139

A

access control lists
 description 91
 examples 100
 ICMP filter example 103
 verifying settings 96
access levels 44
access lists
 adding 95
 configuration commands (table) 97
 deleting 96
 permit-established example 100
 permit-established keyword 95
access masks
 adding 95
 deleting 96
access policies, description 91
access profiles
 reverse mask 105
accounts
 creating 46
 deleting 47
 viewing 46
adding
 access lists 95
 access masks 95
 rate limits 95
Address Resolution Protocol. *See* ARP
admin account 45
Advanced Edge functionality 36
aging entries, FDB 87
alarm actions 157
Alarms, RMON 156
area 0, OSPF 188
areas, OSPF 187
ARP
 clearing entries 178
 communicating with devices outside subnet 171
 configuring proxy ARP 170
 incapable device 170
 proxy ARP between subnets 171
 proxy ARP, description of 170
 responding to ARP requests 170

 table, displaying 172
autonegotiation 71

B

backbone area, OSPF 188
blackhole entries, FDB 88
boot option commands (table) 223
BOOTP
 and UDP-Forwarding 180
BOOTP relay
 configuring 179
BOOTP, using 50
BootROM
 menu, accessing 223
 prompt 223
 upgrading 222
BPDU tunneling 160

C

cable types and distances 22
certification marks 214
CLI
 command history 42
 command shortcuts 40
 line-editing keys 41
 named components 41
 numerical ranges, Summit24e3 switch 40
 symbols 41
 syntax helper 40
 using
 command 42
 history 40
 shortcuts 39
 syntax, understanding
Command-Line Interface. *See* CLI
common commands (table) 42
communicating with devices outside subnet 171
complete configuration download 221
configuration
 downloading 221
 downloading complete 221
 downloading incremental 221
 logging 153
 primary and secondary 219
 saving changes 219

schedule download	222
uploading to file	219
console connection	50
console port	
connecting equipment to	27
controlling Telnet access	52
conventions	
notice icons, About This Guide	17
text, About This Guide	18
creating	
access lists	95
access masks	95
rate limits	95

D

database applications, and QoS	134
database overflow, OSPF	187
default	
passwords	45
settings	38
users	45
default STP domain	160
default VLAN	84
delete	
access list	96
access masks	96
rate limit	96
deleting a session	52
DHCP and UDP-Forwarding	180
DHCP relay, configuring	179
DHCP server	66
DiffServ, configuring	140
dimensions, Summit24e3 switch	213
disabling a switch port	71
disabling route advertising (RIP)	185
disconnecting a Telnet session	52
distance-vector protocol, description	184
DLCS	
configuration commands (table)	145
description	144
guidelines	145
limitations	145
DNS	
configuration commands (table)	47
description	47
Domain Name Service. <i>See</i> DNS	
domains, Spanning Tree Protocol	159
downloading incremental configuration	221
dynamic entries, FDB	87
dynamic routes	169

E

EAPS	
commands (table)	123
domain, creating and deleting	124
enabling and disabling a domain	127
enabling and disabling on a switch	128
polling timers, configuring	125
ring port, unconfiguring	128
show eaps display fields (table)	129
status information, displaying	128
switch mode, defining	125
ECMP. <i>See</i> IP route sharing	

EDP	
commands (table)	77
description	76
electromagnetic compatibility	213
enabling a switch port	71
Equal Cost Multi-Path (ECMP) routing. <i>See</i> IP route sharing	
errors, port	149
establishing a Telnet session	50
Events, RMON	156
export restrictions	37
Extreme Discovery Protocol <i>See</i> EDP	
ExtremeWare	
factory defaults	38
features	19, 33

F

FDB	
adding an entry	88
aging entries	87
blackhole entries	88
configuration commands (table)	89
configuring	88
contents	87
creating a permanent entry example	90
displaying	90
dynamic entries	87
entries	87
non-aging entries	87
permanent entries	87
QoS profile association	88
feature licensing	
Advanced Edge functionality	36
description	36
Edge functionality	36
license keys	37
ordering	37
verifying	36
file server applications, and QoS	135
flow control	72
Forwarding Database. <i>See</i> FDB	
free-standing installation	27
full-duplex	22

G

Greenwich Mean Time Offsets (table)	68
-------------------------------------	----

H

hardware address	21
heat dissipation	213
History, RMON	156

I

ICMP configuration commands (table)	175
IEEE 802.1Q	82
IGMP	
configuration commands (table)	207
description	205
snooping	206
image	
downloading	217
primary and secondary	217
upgrading	217

installation			
free-standing	27		
rack	26		
verifying	28		
interfaces, router	168		
Internet Group Management Protocol. <i>See</i> IGMP			
IP address, entering	51		
IP multicast routing			
configuring	206		
description	205		
disabling	208		
IGMP			
configuration commands (table)	207		
description	205		
snooping	206		
reset and disable commands (table)	208		
resetting	208		
settings, displaying	207		
show commands (table)	207		
IP route sharing	170		
IP TOS configuration commands (table)	140		
IP unicast routing			
basic IP commands (table)	173		
BOOTP relay	179		
configuration examples	177		
configuring	172		
default gateway	167		
description	35		
DHCP relay	179		
disabling	178		
ECMP			
enabling	172		
IP route sharing	170		
proxy ARP	170		
reset and disable commands (table)	178		
resetting	178		
router interfaces	168		
router show commands (table)	178		
routing table			
configuration commands (table)	174		
dynamic routes	169		
multiple routes	169		
populating	168		
static routes	169		
settings, displaying	178		
verifying the configuration	172		
IRDP	176		
K			
keys			
line-editing	41		
port monitoring	150		
L			
LEDs			
Summit24e3	21		
license keys	37		
licensing			
Advanced Edge functionality	36		
description	36		
Edge functionality	36		
license keys	37		
ordering	37		
verifying	36		
line-editing keys	41		
link-state database	186		
link-state protocol, description	184		
load sharing			
algorithms	73		
configuring	74		
description	73		
load-sharing group, description	73		
master port	74		
verifying the configuration	75		
local logging	152		
log display	152		
logging			
and Telnet	152		
commands (table)	153		
configuration changes	153		
description	151		
fault level	151		
local	152		
message	152		
real-time display	152		
remote	153		
subsystem	151		
timestamp	151		
logging in	29, 46		
M			
MAC address	21		
management access	44		
master port			
load sharing	74		
maximum Telnet session	50		
media types and distances	22		
MIBs	54		
mirroring. <i>See</i> port-mirroring			
monitoring the switch	147		
multiple routes	169		
N			
names, VLANs	84		
NAT			
configuration commands (table)	114		
creating rules	116		
rule matching	116		
timeout commands (table)	116		
Network Address Translation. <i>See</i> NAT			
network login	62		
campus mode	62		
configuration example	63		
configuring	63		
user login	64		
configuration commands (table)	66		
DHCP server	63		
disabling	67		
ISP mode	62, 65		
configuration example	65		
configuring	65		
RADIUS server configuration	63		
settings, displaying	66		
non-aging entries, FDB	87		
Not-So-Stubby_Area. <i>See</i> NSSA			
NSSA. <i>See</i> OSPF			

O

opaque LSAs, OSPF	187
Open Shortest Path First. <i>See</i> OSPF	
opening a Telnet session	50
OSPF	
advantages	184
area 0	188
areas	187
backbone area	188
configuration commands (table)	197
consistency	187
database overflow	187
description	184, 186
disabling	203
display filtering	202
enabling	172
hello interval	198
link type	190
link-state database	186
normal area	189
NSSA	188
opaque LSAs	187
point-to-point links	190
redistributing routes	191
reset and disable commands (table)	203
resetting	203
router types	187
routing access policies	107
settings, displaying	202
show commands (table)	202
stub area	188
virtual link	189
wait interval, configuring	201

P

passwords	
default	45
forgetting	46
permanent entries, FDB	87
permit-established keyword	95
ping command	48
poison reverse	185
port	
autonegotiation	71
configuring on Summit24e3 switch	71
connections	22
enabling and disabling	71
errors, viewing	149
monitoring display keys	150
priority, STP	163
receive errors	150
statistics, viewing	149
STP state, displaying	165
STPD membership	160
Summit24e3 switch	71
switch commands (table)	72
transmit errors	149
port-based VLANs	80
port-mirroring	
and protocol analyzers	76
description	75
example	76
switch configuration commands (table)	76
power supply specifications	213

powering on the switch	28
power-off specifications	213
primary image	217
private community, SNMP	54
protocol analyzers, use with port-mirroring	76
proxy ARP	
communicating with devices outside subnet	171
conditions	170
configuring	170
MAC address in response	171
responding to requests	170
subnets	171
table, displaying	178
proxy ARP, description	170
public community, SNMP	54

Q

QoS	
802.1p configuration commands (table)	139
802.1p priority	138
applications	134
blackhole	137
configuration commands (table)	135
database applications	134
description	34, 133
DiffServ, configuring	140
examples	
MAC address	137
source port	142
VLAN	142
FDB entry association	88
file server applications	135
IP TOS configuration commands (table)	140
traffic groupings	136
access list	136
blackhole	137
explicit packet marking	138
MAC address	137
source port	142
VLAN	142
traffic groupings (table)	136
verifying	143
video applications	134
voice applications	134
web browsing applications	135
QoS monitor	
description	143
real-time display	143
Quality of Service. <i>See</i> QoS	133

R

rack mounting the switch	26
RADIUS	
and TACACS+	56, 61
client configuration	56
configuration commands (table)	57
description	56
Merit server configuration (example)	58
per-command authentication	56
per-command configuration (example)	59
RFC 2138 attributes	58
servers	56
TCP port	56


rate limits		
adding	95	
deleting	96	
rate-limiting	144	
receive errors	150	
remote logging	153	
Remote Monitoring. <i>See</i> RMON		
renaming a VLAN	85	
reset button	20	
reset to factory defaults	219	
responding to ARP requests	170	
reverse mask	105	
RIP		
advantages	184	
configuration commands (table)	192	
configuration example	195	
description	184	
disabling route advertising	185	
enabling	172	
limitations	184	
poison reverse	185	
redistributing routes	191	
reset and disable commands (table)	196	
routing access policies	105	
routing table entries	185	
settings, displaying	196	
show commands (table)	196	
split horizon	185	
triggered updates	185	
version 2	185	
RMON		
alarm actions	157	
Alarms group	156	
Events group	156	
features supported	155	
History group	156	
probe	155	
Statistics group	156	
route sharing. <i>See</i> IP route sharing		
router interfaces	168	
router types, OSPF	187	
routing access policies		
access profile		
applying	105	
changing	108	
configuring	104	
creating	104	
types	104	
configuration commands (table)	109	
deny	104	
examples		
OSPF	107	
RIP	106	
none	104	
OSPF	107	
permit	104	
removing	108	
RIP	105	
using	104	
Routing Information Protocol. <i>See</i> RIP		
routing table, populating	168	
routing. <i>See</i> IP unicast routing		
S		
safety information	209	
saving configuration changes	219	
scheduling configuration download	222	
secondary image	217	
security licensing		
description	37	
obtaining	37	
serial port. <i>See</i> console port		
sessions, deleting	52	
shortcuts, command	40	
Simple Network Management Protocol. <i>See</i> SNMP		
size, Summit24e3 switch	213	
SNMP		
community strings	54	
configuration commands (table)	55	
configuring	54	
settings, displaying	55	
supported MIBs	54	
system contact	54	
system location	54	
system name	54	
trap receivers	54	
using	54	
SNTP		
configuration commands (table)	69	
configuring	67	
Daylight Savings Time	67	
description	67	
example	70	
Greenwich Mean Time offset	67	
Greenwich Mean Time Offsets (table)	68	
NTP servers	67	
socket, power	21	
software licensing		
security features	37	
SSH2 protocol	37	
Spanning Tree Protocol. <i>See</i> STP		
speed, ports	71	
split horizon	185	
SSH2 protocol		
authentication key	53	
description	37, 53	
enabling	53	
TCP port number	53	
stand-alone switch		
enabling and disabling ports	71	
static routes	169	
statistics		
port	149	
Statistics, RMON	156	
status monitoring	147	
status monitoring commands (table)	147	
STP		
and VLANs	160	
BPDU tunneling	160	
bridge priority	163	
configurable parameters	163	
configuration commands (table)	163	
configuration example	165	
configuring	162	
default domain	160	
description	34	
disable and reset commands (table)	165	

displaying settings	165	controlling access	52
domains	159	disconnecting a session	52
examples	160	logging	152
forward delay	163	maximum sessions	50
hello time	163	opening a session	50
max age	163	using	50
overview	159	Terminal Access Controller Access Control System Plus. <i>See</i>	
path cost	163	TACACS+	
port priority	163	TFTP	
port state, displaying	165	server	217
stub area, OSPF	188	using	219
Summit24e3 switch		traceroute command	48
certification marks	214	traffic groupings	136
dimensions	213	traffic rate-limiting	144
electromagnetic compatibility	213	transmit errors	149
environmental requirements	214	triggered updates	185
front view	20	trunks	82
heat dissipation	213		
LEDs	21	U	
load sharing	74	UDP-Forwarding	
load sharing example	75	and BOOTP	180
MAC address	21	and DHCP	180
media distances, supported	22	configuration commands (table)	181
media types, supported	22	configuring	180
port configuration	71	description	180
power socket	21	example	181
power supply specifications	213	profiles	180
power-off specifications	213	VLANs	180
rear view	21	upgrading the image	217
reset button	20	uploading the configuration	219
serial number	21	user account	45
size	213	users	
verifying load sharing	75	access levels	44
weight	213	authenticating	56
Summit24e3 switch, general		creating	46
free-standing installation	27	default	45
installing	26	viewing	46
location	25		
port connections	22	V	
powering on	28	verifying the installation	28
rack mounting	26	video applications, and QoS	134
stacking with other devices	27	viewing accounts	46
verifying the installation	28	Virtual LANs. <i>See</i> VLANs	
switch		virtual link, OSPF	189
logging	151	VLAN tagging	82
monitoring	147	VLANs	
RMON features	155	and STP	160
switch port commands (table)	72	assigning a tag	82
syntax, understanding	39	benefits	79
syslog host	153	configuration commands (table)	85
system contact, SNMP	54	configuration examples	86
system location, SNMP	54	configuring	85
system name, SNMP	54	<i>default</i>	84
		description	34
T		disabling route advertising	185
TACACS+		displaying settings	86
and RADIUS	56, 61	mixing port-based and tagged	84
configuration commands (table)	61	names	84
description	60	port-based	80
servers, specifying	61	renaming	85
tagging, VLAN	82	routing	172
technical support	229	tagged	82
Telnet		trunks	82
connecting to another host	50		

types	80
UDP-Forwarding	180
voice applications, QoS	134

W

web browsing applications, and QoS	135
weight, Summit24e3 switch	213



Index of Commands

C

clear counters	153	config iparp timeout	173
clear dlcs	145	config iproute add	174
clear fdb	89, 137	config iproute add blackhole	174
clear igmp snooping	208	config iproute add default	52, 172, 174
clear iparp	173, 178	config iproute delete	174
clear ipfdb	173, 178	config iproute delete blackhole	174
clear ipmc cache	208	config iproute delete default	174
clear log	153	config iproute priority	172, 174
clear session	42, 52	config irdp	175
config access-profile	109	config log display	152, 154
config access-profile add	104, 109	config mirroring add	76
config access-profile delete	105, 109	config mirroring delete	76
config access-profile mode	104	config nat finrst-timeout	116
config account	42	config nat icmp-timeout	116
config banner	42	config nat syn-timeout	117
config bootprelay add	173, 180	config nat tcp-timeout	117
config bootprelay delete	173, 180	config nat timeout	117
config dns-client add	47	config nat udp-timeout	117
config dns-client default-domain	47	config nat vlan	112, 114
config dns-client delete	47	config ospf area nssa	188
config download server	222, 223	config ospf ase-limit	187
config eaps name	124	config ospf add virtual-link	198
config eaps add control vlan	123	config ospf add vlan	198
config eaps add protect vlan	124, 127	config ospf add vlan area link-type	197
config eaps delete control vlan	123	config ospf area add range	198
config eaps delete protect vlan	124	config ospf area delete range	198
config eaps failtime	125	config ospf area external-filter	107, 109
config eaps hellotime	123, 125	config ospf area interarea-filter	107, 109
config eaps mode	123, 125	config ospf area normal	198
config eaps primary port	123, 126	config ospf area nssa	198
config eaps secondary port	123, 126	config ospf area stub	198
config fdb agingtime	89	config ospf asbr-filter	107, 110, 198
config igmp query_interval	207	config ospf ase-limit	198
config igmp snooping	207	config ospf ase-summary add	198
config iparp add	173	config ospf ase-summary delete	199
config iparp add proxy	170, 173	config ospf authentication	197
config iparp delete	173	config ospf cost	197
config iparp delete proxy	173	config ospf delete virtual-link	199
		config ospf delete vlan	199

config ospf direct-filter	107, 110, 199	config tacacs-accounting	61
config ospf lsa-batching-timer	199	config tacacs-accounting shared-secret	61
config ospf metric-table	199	config time	43
config ospf originate-default	201	config timezone	43, 67
config ospf routerid	199	config udp-profile add	181
config ospf spf-hold-time	199	config udp-profile delete	181
config ospf timer	198	config vlan add port	85
config ospf vlan	199	config vlan delete port	85
config ospf vlan area	188	config vlan dhcp-address-range	66
config ospf vlan neighbor add	197	config vlan dhcp-lease-timer	66
config ospf vlan neighbor delete	197	config vlan dhcp-options	66
config ospf vlan timer	200, 201	config vlan ipaddress	43, 52, 85, 172
config ports auto off	42, 71, 72	config vlan name	85, 86
config ports auto on	72	config vlan netlogin-lease-timer	66
config ports display-string	72	config vlan priority	139
config ports qosprofile	135, 142	config vlan qosprofile	135, 142
config radius server	57	config vlan tag	85
config radius shared-secret	57	config vlan udp-profile	182
config radius-accounting	57	configure eaps failtime	123
config radius-accounting shared-secret	57	create access-list	95, 97
config rip add	192	create access-mask	95, 98
config rip delete	193, 196	create access-profile type	104, 110
config rip garbagetime	193	create account	43, 46
config rip routetimeout	193	create eaps	124
config rip rxmode	193	create fdbentry	89, 137
config rip txmode	193	create fdbentry blackhole	89
config rip updatetime	193	create fdbentry dynamic	89
config rip vlan cost	193	create ospf area	188, 200
config rip vlan export-filter	106, 110	create rate-limit	95, 99
config rip vlan import-filter	105, 110	create stpd	162, 164
config rip vlan trusted-gateway	105, 110	create udp-profile	182
config sharing address-based	72, 74	create vlan	43, 86
config snmp add trapreceiver	55		
config snmp community	55	D	
config snmp delete trapreceiver	55	delete access-list	96, 99
config snmp syscontact	55	delete access-mask	96, 100
config snmp syslocation	55	delete access-profile	110
config snmp sysname	55	delete account	43
config snmp-client	67	delete eaps	124
config snmp-client server	69	delete fdbentry	89
config snmp-client update-interval	68, 70	delete ospf area	203
config ssh2 key	42, 53	delete rate-limit	96, 100
config stpd add vlan	162, 163	delete stpd	165
config stpd forwarddelay	163	delete udp-profile	182
config stpd hellotime	163	delete vlan	43, 86
config stpd maxage	164	disable bootp	43, 173, 178
config stpd port cost	164	disable bootprelay	173, 179
config stpd port priority	164	disable cli-config-logging	43, 153, 154
config stpd priority	164	disable clipaging	43
config syslog	153, 154	disable dhcp ports vlan	66
config syslog delete	154	disable diffserv examination ports	140
config sys-recovery-level	43, 151	disable dlcs	145
config tacacs	61	disable dlcs ports	145
config tacacs shared-secret	61	disable eaps	124, 127, 128

disable edp ports	77	disable snmp-client	70
disable icmp	179	disable ssh2	44
disable icmp address-mask	179	disable stpd	165
disable icmp parameter-problem	175	disable stpd port	165
disable icmp port-unreachables	179	disable syslog	154
disable icmp redirects	179	disable tacacs	61
disable icmp time-exceeded	179	disable tacacs-accounting	61
disable icmp timestamp	179	disable tacacs-authorization	61
disable icmp unreachable	179	disable telnet	44, 52
disable icmp userredirects	179	download bootrom	47, 223
disable idletimeouts	43	download configuration	47, 221, 223
disable igmp	208	download configuration cancel	222, 223
disable igmp snooping	208	download configuration every	222, 223
disable ignore-bpdu	160	download configuration incremental	222
disable ignore-bpdu vlan	165	download image	47, 217, 224
disable ignore-stp vlan	165		
disable ipforwarding	173, 179	E	
disable ipforwarding broadcast	173, 179	enable bootp	44, 173
disable ipforwarding fast-direct-broadcast	170	enable bootp vlan	51
disable ipforwarding ignore-broadcast	170	enable bootprelay	173, 179
disable ipmcf forwarding	208	enable cli-config-logging	44, 153, 154
disable ip-option loose-source-route	175	enable clipaging	44
disable ip-option record-route	175	enable dhcp ports	66
disable ip-option record-timestamp	175	enable diffserv examination ports	140, 141
disable ip-option strict-source-route	175	enable dlcs	145
disable ip-option use-router-alert	175	enable dlcs ports	145
disable iproute sharing	174	enable eaps	124, 127, 128
disable irdp	179	enable edp ports	77
disable learning port	89	enable icmp address-mask	175
disable log display	154	enable icmp parameter-problem	176
disable loopback-mode vlan	173	enable icmp port-unreachables	176
disable mirroring	76	enable icmp redirects	176
disable nat	117	enable icmp time-exceeded	176
disable netlogin ports vlan	66	enable icmp timestamp	176
disable ospf	203	enable icmp unreachable	176
disable ospf capability opaque-lsa	187, 200	enable icmp userredirects	176
disable ospf export	169	enable idletimeouts	44
disable ospf export direct	203	enable igmp	207
disable ospf export rip	192, 203	enable igmp snooping	207
disable ospf export static	192, 203	enable ignore-bpdu	160
disable ports	43, 71, 72	enable ignore-bpdu vlan	164
disable radius	57	enable ignore-stp vlan	164
disable radius-accounting	57	enable ipforwarding	172, 174
disable rip	196	enable ipforwarding broadcast	174
disable rip aggregation	196	enable ipforwarding fast-direct-broadcast	170
disable rip export	169, 192, 196	enable ipforwarding ignore-broadcast	170
disable rip originate-default	196	enable ipmcf forwarding	206
disable rip poisonreverse	196	enable ip-option loose-source-route	176
disable rip splithorizon	196	enable ip-option record-route	176
disable rip triggerupdates	196	enable ip-option record-timestamp	176
disable rmon	156	enable ip-option strict-source-route	176
disable sharing	72, 75	enable ip-option use-router-alert	176
disable snmp access	55	enable iproute sharing	175
disable snmp traps	55	enable irdp	176

enable learning port	89
enable log display	152, 154
enable loopback-mode vlan	174
enable mirroring	76
enable nat	114
enable netlogin ports	66
enable ospf export direct	200
enable ospf	172, 200
enable ospf capability opaque-lsa	187, 200
enable ospf export	169
enable ospf export rip	192, 200
enable ospf export static	192, 201
enable ports	71, 72
enable radius	57
enable radius-accounting	58
enable rip	172, 193
enable rip aggregation	194
enable rip export	169, 192, 194
enable rip originate-default	194
enable rip poisonreverse	194
enable rip splithorizon	194
enable rip triggerupdates	194
enable rmon	156
enable route sharing	170
enable sharing	72, 75
enable snmp access	55
enable snmp traps	55
enable snmp-client	67, 70
enable ssh2	44, 53
enable stpd	163, 164
enable stpd port	164
enable syslog	153, 154
enable tacacs	61
enable tacacs-accounting	61
enable tacacs-authorization	61
enable telnet	44, 53

H

history	42, 44
---------	--------

L

logout	52
--------	----

N

nslookup	47
----------	----

P

ping	47, 48
------	--------

Q

quit	52
------	----

R

reboot	218, 224
restart ports	73
rtlookup	175

S

save	52, 219, 224
show access-list	96, 100
show access-mask	96, 100
show access-profile	110
show accounts	46
show banner	44
show configuration	224
show debug-tracing	229
show diagnostics	147
show dlcs	145
show dns-client	47
show eaps	124, 128
show edp	77
show fdb	90
show fdb permanent	137, 144
show igmp snooping	207
show iparp	172, 178
show iparp proxy	178
show ipconfig	172, 178, 180
show ipfdb	172, 178
show ipmc cache	208
show iproute	172, 178
show ipstats	178
show log	148, 152, 154
show log config	148, 155
show management	52, 55
show memory	148
show mirroring	76
show nat connections	117
show nat rules	117
show nat stats	117
show nat timeout	117
show ospf	192, 202
show ospf area	202
show ospf ase-summary	202
show ospf interfaces	202
show ospf lsdb	202
show ospf virtual-link	202
show ports collisions	73
show ports configuration	73, 75
show ports info	73, 141, 142, 144
show ports packet	73
show ports qosmonitor	143
show ports rxerrors	73, 150
show ports stats	73, 149
show ports txerrors	73, 149
show ports utilization	73
show qosprofile	137, 142, 143

show radius	58
show radius-accounting	58
show rate-limit	96, 100
show rip	196
show rip stat	196
show rip vlan	196
show session	52
show sharing address-based	73, 74
show snmp client	68
show snmp-client	70
show stpd	165
show stpd port	165
show switch	68, 144, 148, 222
show tacacs	62
show tacacs-accounting	62
show tech-support	148
show udp-profile	182
show version	148
show vlan	86, 142, 144

T

telnet	47, 50
traceroute	47, 48

U

unconfig eaps	124
unconfig eaps primary port	128
unconfig eaps secondary port	128
unconfig icmp	176, 179
unconfig igmp	208
unconfig irdp	177, 179
unconfig management	55
unconfig ospf	203
unconfig ports display-string	73
unconfig ports monitor vlan	86
unconfig radius	58
unconfig radius-accounting	58
unconfig rip	196
unconfig stpd	166
unconfig switch	44, 219
unconfig switch all	219
unconfig tacacs	62
unconfig tacacs-accounting	62
unconfig udp-profile	182
unconfig vlan ipaddress	86
upload configuration	47, 220, 224
upload configuration cancel	220, 224
use configuration	219, 224
use image	224

