ExtremeWare® Software Command Reference Guide

Software Version 6.2.2

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

Published: June 2002
Part number: 100115-00 rev 03

Authors: Richard Small, Valerie Swisher, Julie Laccabue

Editor: Richard Small, Julie Laccabue

Production: Julie Laccabue

Special Thanks: Theresa Zamborsky

Contents

Chapter 3 Commands for Managing the Switch

Chapter 18 PoS Commands

Index of Commands

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 complete syntax for all the commands available in the currently-supported versions of the ExtremeWare® software running on either modular or stand-alone switches from Extreme Networks. This also includes commands that support specific modules such as the ARM, MPLS or PoS modules.

This guide is intended for use as a reference by network administrators who are responsible for installing and setting up network equipment. It assumes knowledge of Extreme switch configuration. For conceptual information and guidance on configuring Extreme switches, refer to the ExtremeWare Software User Guide for your version of the ExtremeWare software.

Terminology

When features, functionality, or operation is specific to a modular or stand-alone switch family, the family name is used. Explanations about features and operations that are the same across all product families simply refer to the product as the “switch.”

Conventions

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

Table 1: Notice Icons

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

Table 1: Notice Icons

Icon	Notice Type	Alerts you to...
	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].
Letter in **bold** type	Letters within a command that appear in bold type indicate the keyboard shortcut for a command. When entering the command, you can use just the bolded letters instead of the entire word.
Words in *italicized* type	Italics emphasize a point or denote new terms at the place where they are defined in the text.

Related Publications

The publications related to this one are:

- *ExtremeWare release notes.*
- *ExtremeWare Software User Guide*
- *ExtremeWare 6.2.1 Software Quick Reference Guide*
- *Extreme Networks Consolidated Hardware Guide*

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

http://www.extremenetworks.com/

1 Command Reference Overview

Introduction

This guide provides details of the command syntax for all ExtremeWare commands as of ExtremeWare version 6.2.2. This includes commands specific to stand-alone and modular switches based on the “i” series chipset. Commands that are supported in ExtremeWare IP Services Technology Releases for specific modules, such as the ARM, PoS, and MPLS modules are also included.

For historical reasons, commands and command options supported in ExtremeWare 4.1 on switches based on the original Summit chipset are also included. This guide does not cover the Summit e-series switches.

This guide does not provide feature descriptions, explanations of the technologies, or configuration examples. For information about the various features and technologies supported by Extreme switches, see the installation and user guides for your product. This guide does not replace the installation and user guides; this guide supplements the installation and user guides.

Audience

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 the following:

- Local area networks (LANs)
- Ethernet concepts
- Ethernet switching and bridging concepts
- Routing concepts
- Internet Protocol (IP) concepts
- Routing Information Protocol (RIP) and Open Shortest Path First (OSPF) concepts
- Border Gateway Protocol (BGP-4) concepts
- IP Multicast concepts
- Distance Vector Multicast Routing Protocol (DVMRP) concepts
- Protocol Independent Multicast (PIM) concepts
- Internet Packet Exchange (IPX) concepts

- Server Load Balancing (SLB) concepts
- Simple Network Management Protocol (SNMP)

This guide also assumes that you have read the Installation and User Guide for your product.

Structure of this Guide

This guide documents each ExtremeWare command. Related commands are grouped together and organized into chapters based on their most common usage. The chapters reflect the organization of the *ExtremeWare Software User Guide*. If a specific command is relevant to a wide variety of functions and could be included in a number of different chapters, we have attempted to place the command in the most logical chapter. Within each chapter, commands appear in alphabetical order. You can use the Index of Commands to locate specific commands if they do not appear where you expect to find them.

For each command, the following information is provided:

- Command Syntax: The actual syntax of the command. The syntax conventions (the use of braces or curly brackets, for example) are defined in the section "Understanding the Command Syntax" on page 41.
- Description: A brief (one sentence) summary of what the command does.
- Syntax Description: The definition of any keywords and options used in the command.
- Default: The defaults, if any, for this command. The default can be the default action of the command if optional arguments are not provided, or it can be the default state of the switch (such as for an enable / disable command).
- Usage Guidelines: Information to help you use the command. This may include prerequisites, prohibitions, and related commands, as well as other information.
- Example: Examples of the command usage, including output, if relevant.
- History: The version of ExtremeWare in which the command was introduced, and version(s) where it was modified, if appropriate.
- Platform Availability: The platforms on which the command is supported. For example, many commands are supported only on platforms based on the "i" series chipset. Other commands may be unique to a particular module. The Summt e-series switches are not included.

NOTE

Commands designated as "available on all platforms" are supported on both Summit chipset-based and "i"-series platforms. Summit e-series devices are not included.

Understanding the Command Syntax

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.

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 3 summarizes command syntax symbols.

Table 3: Command Syntax Symbols

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

Command Completion with Syntax Helper

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

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

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

Abbreviated Syntax

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

NOTE

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

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.

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, instead of entering the modular switch command

```
config vlan engineering delete port 1:3,4:6
```

you could enter the following shortcut:

```
config engineering delete port 1:3,4:6
```

Similarly, 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
```

Modular Switch Numerical Ranges

Commands that require you to enter one or more port numbers on a modular switch use the parameter `<portlist>` in the syntax. A `<portlist>` can be one port on a particular slot. For example,

```
port 3:1
```

A `<portlist>` can be a range of numbers. For example,

```
port 3:1-3:3
```

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

```
port 3:1,4:8,6:10
```

You can specify all ports on a particular slot. For example,

```
port 3:*
```

indicates all ports on slot 3.

You can specify a range of slots and ports. For example,

```
port 2:3-4:5
```

indicates slot 2, port 3 through slot 4, port 5.

Stand-alone Switch Numerical Ranges

Commands that require you to enter one or more port numbers on a stand-alone 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
```

Line-Editing Keys

Table 4 describes the line-editing keys available using the CLI

Table 4: Line-Editing Keys

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

Command History

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

```
history
```

2 Commands for Accessing the Switch

This chapter describes:

- Commands used for accessing and configuring the switch including how to set up user accounts, passwords, date and time settings, and software licenses
- Commands used for configuring the Domain Name Service (DNS) client
- Commands used for checking basic switch connectivity

ExtremeWare supports the following two levels of management:

- User
- Administrator

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.

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.

The DNS client in ExtremeWare augments certain ExtremeWare commands to accept either IP addresses or host names. For example, DNS can be used during a Telnet session when you are accessing a device or when using the `ping` command to check the connectivity of a device.

The switch offers the following commands for checking basic connectivity:

- `ping`
- `traceroute`

The `ping` command enables you to send Internet Control Message Protocol (ICMP) echo messages to a remote IP device. The `traceroute` command enables you to trace the routed path between the switch and a destination endstation.

clear session

```
clear session <number>
```

Description

Terminates a Telnet session from the switch.

Syntax Description

number	Specifies a session number from `show session` output to terminate.

Default

N/A.

Usage Guidelines

An administrator-level account can disconnect a management session that has been established by way of a Telnet connection. You can determine the session number of the session you want to terminate by using the `show session` command. The `show session` output displays information about current Telnet sessions including:

- The session number
- The login date and time
- The user name
- The type of Telnet session

Depending on the software version running on your switch, additional session information may be displayed. The session number is the first number displayed in the `show session` output.

Example

The following command terminates session 4 from the system:

```
clear session 4
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config account

```
config account <user account> {encrypted} {<password>}
```

Description

Configures a user account password.

Syntax Description

user account	Specifies a user account name.
encrypted	This option is for use only by the switch when generating an ASCII configuration file. Specifies that the password should be encrypted when the configuration is uploaded to a file. Should not be used through the CLI.
password	Specifies a user password. Supported in versions 4.x and 6.0.x only. In versions 6.1 and later, the switch will prompt for entry of the password interactively.

Default

N/A.

Usage Guidelines

You must create a user account before you can configure a user account. Use the `create account` command to create a user account.

You must have administrator privileges to change passwords for accounts other than your own. User names and passwords are case-sensitive.

The `encrypted` option is used by the switch when generating an ASCII configuration file (using the `upload configuration` command), and parsing a switch-generated configuration file (using the `download configuration` command). Do not select the encrypted option in the CLI.

For version 6.1 and higher:

- The password cannot be specified on the command line. Instead, the switch will interactively prompt you to enter the password, and will then prompt you to reenter the password to verify that you have entered it correctly.

For version 6.0 and higher:

- Passwords must have a minimum of 1 character and can have a maximum of 32 characters.

For version 4.x:

- Passwords must have a minimum of 4 characters and can have a maximum of 12 characters.

Example

The following command defines a new password for the account *admin*:

```
config account admin
```

The switch responds with a password prompt:

```
password:
```

Your keystrokes will not be echoed as you enter the new password. After you enter the password, the switch will then prompt you to reenter it.

```
Reenter password:
```

Assuming you enter it successfully a second time, the password is now changed.

In ExtremeWare version 4.1.19, the following command defines a new password, *Extreme1*, for the account *admin*:

```
config account admin Extreme1
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config banner

```
config banner
```

Description

Configures the banner string that is displayed at the beginning of each login prompt of each session.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

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.

For version 6.0 and higher:

- You can enter up to 24 rows of 79-column text that is displayed before the login prompt of each session.

For version 2.0 and 4.x:

- You can enter up to 24 rows of 80-column text that is displayed before the login prompt of each session.

Example

The following command adds a banner, *Welcome to the switch*, before the login prompt:

```
config banner [Return]
Welcome to the switch
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config banner netlogin

```
config banner netlogin
```

Description

Configures the network login banner that is displayed at the beginning of each login prompt of each session.

Syntax Description

This command has no arguments or variables.

Default

N / A.

Usage Guidelines

The network login banner and the switch banner cannot be used at the same time. If you configure a Network Login banner, users do **not** see the normal banner. If no banner is configured, the Extreme logo is displayed. The network login banner displays in HTML. No links or images are supported.

Press [Enter] to enter text on a new line. Press [Enter] twice to finish entering the network login banner. You can enter up to 1024 characters in the banner.

Example

The following command adds the banner “Welcome to your switch” in 8 point purple Arial before the login prompt:

```
config banner netlogin [Enter]
<font face="Arial" size=8 color=534579></font>Welcome to your switch
[Enter]
[Enter]
```

History

This command was introduced in ExtremeWare 6.2.2.

Platform Availability

This command is available on all “*i*” series platforms.

config dns-client add

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

Description

Adds a DNS name server to the available server list for the DNS client.

Syntax Description

ipaddress	Specifies an IP address.

Default

N / A.

Usage Guidelines

Up to three DNS name servers can be configured in ExtremeWare versions prior to 6.2.1. In ExtremeWare 6.2.1 and later, eight DNS name servers can be configured.

Example

The following command specifies that the switch use the DNS server 10.1.2.1:

```
config dns-client add 10.1.2.1
```

History

This command was first available in ExtremeWare 4.0.

This command was modified in ExtremeWare 6.2.1 to support up to eight DNS name servers.

Platform Availability

This command is available on all platforms.

config dns-client add domain-suffix

```
config dns-client add domain-suffix <domain_name>
```

Description

Adds a domain name to the domain suffix list.

Syntax Description

domain_name	Specifies a domain name.

Default

N/A.

Usage Guidelines

The domain suffix list can include up to 6 items. The most recently added entry on the domain suffix list will be the last name used during name resolution if the use of all previous names fails to resolve a name. This command will not overwrite any exiting entries. If a null string is used as the last suffix in the list, and all other lookups fail, the name resolver will attempt to look up the name with no suffix.

Example

The following command configures a domain name and adds it to the domain suffix list:

```
config dns-client add domain-suffix xyz_inc.com
```

History

This command was first available in ExtremeWare 6.2.2.

Platform Availability

This command is available on all platforms.

config dns-client add name-server

```
config dns-client add name-server <ipaddress>
```

Description

Adds a DNS name server to the available server list for the DNS client.

Syntax Description

ipaddress	Specifies an IP address.

Default

N / A.

Usage Guidelines

Up to three DNS name servers can be configured in ExtremeWare versions prior to 6.2.1. In ExtremeWare 6.2.1 and later, eight DNS name servers can be configured.

Example

The following command specifies that the switch use the DNS server 10.1.2.1:

```
config dns-client add 10.1.2.1
```

History

This command was first available in ExtremeWare 6.2.2.

Platform Availability

This command is available on all platforms.

config dns-client default-domain

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

Description

Configures the domain that the DNS client uses if a fully qualified domain name is not entered.

Syntax Description

domain_name	Specifies a default domain name.

Default

N/A.

Usage Guidelines

Sets the DNS client default domain name to `domain_name`. The default domain name will be used to create a fully qualified host name when a domain name is not specified. For example, if the default default domain name is set to "food.com" then when a command like "`ping dog`" is entered, the ping will actually be executed as "`ping dog.food.com`".

Example

The following command configures the default domain name for the server:

```
config dns-client default-domain xyz_inc.com
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

config dns-client delete

```
config dns-client delete <ipaddress>
```

Description

Removes a DNS name server from the available server list for the DNS client.

Syntax Description

ipaddress	Specifies an IP address.

Default

N/A.

Usage Guidelines

None

Example

The following command removes a DNS server from the list:

```
config dns-client delete 10.1.2.1
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

config dns-client delete domain-suffix

```
config dns-client delete domain-suffix <domain_name>
```

Description

Deletes a domain name from the domain suffix list.

Syntax Description

domain_name	Specifies a domain name.

Default

N/A.

Usage Guidelines

This command randomly removes an entry from the domain suffix list. If the deleted item was not the last entry in the list, all items that had been added later are moved up in the list. If no entries in the list match the domain name specified, an error message will be displayed.

Example

The following command deletes a domain name from the domain suffix list:

```
config dns-client delete domain-suffix xyz_inc.com
```

History

This command was first available in ExtremeWare 6.2.2.

Platform Availability

This command is available on all platforms.

config dns-client delete name-server

```
config dns-client delete name-server <ipaddress>
```

Description

Removes a DNS name server from the available server list for the DNS client.

Syntax Description

ipaddress	Specifies an IP address.

Default

N/A.

Usage Guidelines

None.

Example

The following command removes a DNS server from the list:

```
config dns-client delete 10.1.2.1
```

History

This command was first available in ExtremeWare 6.2.2.

Platform Availability

This command is available on all platforms.

config idletimeouts

```
config idletimeouts <minutes>
```

Description

Configures the time-out for idle HTTP, console, and Telnet sessions.

Syntax Description

minutes	Specifies the time-out interval, in minutes. Range is 1 to 240 (1 minute to 4 hours).

Default

Default time-out is 20 minutes.

Usage Guidelines

This command configures the length of time the switch will wait before disconnecting idle HTTP, console, or Telnet sessions. The idletimeouts feature must be enabled for this command to have an effect (the idletimeouts feature is disabled by default).

In ExtremeWare v 6.2.0, the time-out interval was specified in seconds, not minutes.

Example

The following command sets the time-out for idle HTTP, login and console sessions to ten minutes:

```
config idletimeouts 10
```

History

This command was first available in ExtremeWare 6.2

This command was modified in ExtremeWare 6.2.1 to change the time-out value specification to minutes.

Platform Availability

This command is available on the "*i*" series switches.

config msm-failover link-action

```
config msm-failover link-action [keep-links-up | take-links-down]
```

Description

Syntax Description

keep-links-up	Configures the external ports to not be reset when MSM failover occurs. This option is available on the "*i*" series switches only.
take-links-down	Configures the external ports to be reset when MSM failover occurs. This option is available on the "*i*" series switches only.

Default

Take-links-down.

Usage Guidelines

When MSM failover occurs, external ports will not be reset if the `keep-links-up` option is configured. When the `keep-links-up` option is configured, peer connections will not notice a link-down indication.

The `keep-links-up` and `take-links-down` options are available on the "*i*" series switches only.

Example

The following command prevents external ports from being reset when an MSM failover occurs:

```
config msm-failover link-action keep-links-up
```

History

This command was first available in ExtremeWare 6.2.2

Platform Availability

This command is available on the BlackDiamond only.

config sys-recovery-level

```
config sys-recovery-level [none | [critical | all] [shutdown | reboot |
msm-failover]]
```

Description

Configures a recovery option for instances where an exception occurs in ExtremeWare.

Syntax Description

none	Configures the level to no recovery. No action is taken when a task exception occurs; there is no system shutdown or reboot.
critical	Configures ExtremeWare to log an error into the syslog and either shutdown or reboot the system after a critical task exception.
all	Configures ExtremeWare to log an error into the syslog and either shutdown or reboot the system after any task exception occurs.
shutdown	Shuts down the switch.
reboot	Reboots the switch.
msm-failover	Triggers the slave MSM64i to take over control of the switch if there is a software exception on the master MSM64i. BlackDiamond only.

Default

None.

Usage Guidelines

This command is used for system troubleshooting. If the system fails before the switch is booted up, the switch will automatically start the console and allow access to the system to view the logs or debug the failure. You can also configure the system to respond to software failures automatically. You must specify one of the following parameters for the system to respond to software failures:

- none—No action is taken when a task exception occurs.
- `critical`—The system will reboot or shutdown if a critical task exception occurs. Critical tasks include the tBGTask, tNetTask, and tESRPTask.
- `all`—The system will reboot or shut down if any task exception occurs.

For version 6.1, the system will always reboot after a task exception when the system recovery level is specified as `all` or `critical`.

For version 6.2 or later, you must specify whether the system should shut down or reboot upon a task exception if the recovery level is `all` or `critical`.

For version 6.2.2 or later, if `msm-failover` is specified on a BlackDiamond and there is a software exception on the Master MSM64*i*, the interrupt handler triggers the Slave MSM64*i* to take over control of the switch.

Example

The following command configures a switch to reboot after a critical task exception occurs:

```
config sys-recovery-level critical reboot
```

The following command configures the Master MSM64*i* to failover to the Slave MSM64*i* if a software exception occurs:

```
config sys-recovery-level critical msm-failover
```

History

This command was first available in ExtremeWare 6.1.

Modified in ExtremeWare 6.2 to support the `shutdown` and `reboot` options.

Modified in ExtremeWare 6.2.2 to support the `msm-failover` option.

Platform Availability

This command is available on all *i*-series switches. The `msm-failover` option is available on BlackDiamond only.

config time

```
config time <date> <time>
```

Description

Configures the system date and time.

Syntax Description

date	Specifies the date in `mm/dd/yyyy` format.
time	Specifies the time in `hh:mm:ss` format.

Default

N/A.

Usage Guidelines

The format for the system date and time is as follows:

```
mm/dd/yyyy hh:mm:ss
```

The time uses a 24-hour clock format. The AM hours range from 1 through 11, and the PM hours range from 12 through 23.

For version 6.0 and higher:

- You cannot set the year past 2036.

For version 2.0 and 4.x:

- You cannot set the year past 2023.

Example

The following command configures a system date of February 15, 2002 and a system time of 8:42 AM and 55 seconds:

```
config time 02/15/2002 08:42:55
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config timezone

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

Description

Configures the Greenwich Mean Time (GMT) offset and Daylight Saving Time (DST) preference.

Syntax Description

GMT_offset	Specifies a Greenwich Mean Time (GMT) offset, in + or - minutes.
std-timezone-ID	Specifies an optional name for this timezone specification. May be up to six characters in length. The default is an empty string.
autodst	Enables automatic Daylight Saving Time.
dst-timezone-ID	Specifies an optional name for this DST specification. May be up to six characters in length. The default is an empty string.
dst_offset	Specifies an offset from standard time, in minutes. Value is in the range of 1 to 60. Default is 60 minutes.
floating_day	Specifies the day, week, and month of the year to begin or end DST each year. Format is: <week><day><month> where: • <week> is specified as [first | second | third | fourth | last] or 1-5 • <day> is specified as [sunday | monday | tuesday | wednesday | thursday | friday | saturday] or 1-7 (where 1 is Sunday) • <month> is specified as [january | february | march | april | may | june | july | august | september | october | november | december] or 1-12 Default for beginning is first sunday april; default for ending is last sunday october.
absolute_day	Specifies a specific day of a specific year on which to begin or end DST. Format is: <month>/<day>/<year> where: • <month> is specified as 1-12 • <day> is specified as 1-31 • <year> is specified as 1970 - 2035 The year must be the same for the begin and end dates.
time_of_day	Specifies the time of day to begin or end Daylight Saving Time. May be specified as an hour (0-23) or as hour:minutes. Default is 2:00.
noautodst	Disables automatic Daylight Saving Time.

Default

Autodst, beginning every first Sunday in April, and ending every last Sunday in October.

Usage Guidelines

Network Time Protocol (NTP) server 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 geographic location.

The `gmt_offset` is specified in +/- minutes from the GMT time.

Automatic DST changes can be enabled or disabled. The default configuration, where DST begins on the first Sunday in April at 2:00 AM and ends the last Sunday in October at 2:00 AM, applies to most of North America, and can be configured with the following syntax:
`configure timezone <gmt_offst> autodst`.

As of ExtremeWare version 6.2.1, the starting and ending date and time for DST may be specified, as these vary in time zones around the world.

- Use the `every` keyword to specify a year-after-year repeating set of dates (e.g. the last Sunday in March every year)
- Use the `on` keyword to specify a non-repeating, specific date for the specified year. If you use this option, you will need to specify the command again every year.
- The `begins` specification defaults to `every first sunday april`.
- The `ends` specification defaults to `every last sunday october`.
- The `ends` date may occur earlier in the year than the `begins` date. This will be the case for countries in the Southern Hemisphere.
- If you specify only the starting or ending time (not both) the one you leave unspecified will be reset to its default.
- The `time_of_day` specification defaults to `2:00`
- The timezone IDs are optional. They are used only in the display of timezone configuration information in the `show switch` command.

To disable automatic DST changes, re-specify the GMT offset using the `noautodst` option:
`configure timezone <gmt_offst> noautodst`.

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 5 describes the GMT offsets.

Table 5: 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

Table 5: Greenwich Mean Time Offsets (continued)

GMT Offset in Hours	GMT Offset in Minutes	Common Time Zone References	Cities
-5:00	-300	EST - Eastern Standard	Bogota, Columbia; Lima, Peru; New York, NY, Trevor City, MI USA
-6:00	-360	CST - Central Standard	Mexico City, Mexico
-7:00	-420	MST - Mountain Standard	Saskatchewan, Canada
-8:00	-480	PST - Pacific Standard	Los Angeles, CA, Cupertino, CA, Seattle, WA USA
-9:00	-540	YST - Yukon Standard	
-10:00	-600	AHST - Alaska-Hawaii Standard	
		CAT - Central Alaska	
		HST - Hawaii Standard	
-11:00	-660	NT - Nome	
-12:00	-720	IDLW - International Date Line West	
+1:00	+60	CET - Central European	Paris, France; Berlin, Germany; Amsterdam, The Netherlands; Brussels, Belgium; Vienna, Austria; Madrid, Spain; Rome, Italy; Bern, Switzerland; Stockholm, Sweden; Oslo, Norway
		FWT - French Winter	
		MET - Middle European	
		MEWT - Middle European Winter	
		SWT - Swedish Winter	
+2:00	+120	EET - Eastern European, Russia Zone 1	Athens, Greece; Helsinki, Finland; Istanbul, Turkey; Jerusalem, Israel; Harare, Zimbabwe
+3:00	+180	BT - Baghdad, Russia Zone 2	Kuwait; Nairobi, Kenya; Riyadh, Saudi Arabia; Moscow, Russia; Tehran, Iran
+4:00	+240	ZP4 - Russia Zone 3	Abu Dhabi, UAE; Muscat; Tblisi; Volgograd; Kabul
+5:00	+300	ZP5 - Russia Zone 4	
+5:30	+330	IST – India Standard Time	New Delhi, Pune, Allahabad, India
+6:00	+360	ZP6 - Russia Zone 5	
+7:00	+420	WAST - West Australian Standard	
+8:00	+480	CCT - China Coast, Russia Zone 7	
+9:00	+540	JST - Japan Standard, Russia Zone 8	
+10:00	+600	EAST - East Australian Standard	
		GST - Guam Standard	
		Russia Zone 9	
+11:00	+660		
+12:00	+720	IDLE - International Date Line East	Wellington, New Zealand; Fiji, Marshall Islands
		NZST - New Zealand Standard	
		NZT - New Zealand	

Example

The following command configures GMT offset for Mexico City, Mexico and disables automatic DST:

```
config timezone -360 noautodst
```

The following four commands are equivalent, and configure the GMT offset and automatic DST adjustment for the US Eastern timezone, with an optional timezone ID of EST:

```
config timezone name EST -300 autodst name EDT 60 begins every first sunday april at
2:00 ends every last sunday october at 2:00

config timezone name EST -300 autodst name EDT 60 begins every 1 1 4 at 2:00 ends
every 5 1 10 at 2:00

config timezone name EST -300 autodst name EDT

config timezone -300 autodst
```

The following command configures the GMT offset and automatic DST adjustment for the Middle European timezone, with the optional timezone ID of MET:

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

The following command configures the GMT offset and automatic DST adjustment for New Zealand. The ending date must be configured each year because it occurs on the first Sunday on or after March 5:

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

History

This command was first available in ExtremeWare 4.0.

Modified in ExtremeWare 6.2.1 to allow configuration of a beginning and ending time for the automatic DST.

Platform Availability

This command is available on all platforms.

create account

```
create account [admin | user] <username> {encrypted} {<password>}
```

Description

Creates a new user account.

Syntax Description

admin	Specifies an access level for account type `admin`.
user	Specifies an access level for account type `user`.
username	Specifies a new user account name. See “Usage Guidelines” for more information.
encrypted	Specifies an encrypted option.
password	Specifies a user password. See “Usage Guidelines” for more information.

Default

By default, the switch is configured with two accounts with the access levels shown in Table 6:

Table 6: User Account Levels

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. This user has access to the `ping` command.

You can use the default names (*admin* and *user*), or you can create new names and passwords for the accounts. Default accounts do not have passwords assigned to them.

Usage Guidelines

The switch can have a total of 16 user accounts. There must be one administrator account on the system.

You must have administrator privileges to change passwords for accounts other than your own. User names and passwords are case-sensitive.

For version 6.0 and higher:

- User account names must have a minimum of 1 character and can have a maximum of 32 characters.
- Passwords must have a minimum of 0 characters and can have a maximum of 16 characters.

For version 4.x and higher:

- Admin-level users and users with RADIUS command authorization can use the `create account` command.

For version 4.x:

- User account name specifications are not available.
- Passwords must have a minimum of 4 characters and can have a maximum of 12 characters.
- The encrypted option should only be used by the switch to generate an ASCII configuration (using the `upload configuration` command), and parsing a switch-generated configuration (using the `download configuration` command).

Example

The following command creates a new account named John2 with administrator privileges:

```
create account admin john2
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 4.0 to support the encrypted option. In addition, admin-level users with RADIUS command authorization were allowed to use the `create account` command.

Platform Availability

This command is available on all platforms.

delete account

```
delete account <username>
```

Description

Deletes a specified user account.

Syntax Description

username	Specifies a user account name.

Default

N/A

Usage Guidelines

Use the `show accounts` command to determine which account you want to delete from the system. The show accounts output displays the following information in a tabular format:

- The user name
- Access information associated with each user
- User login information
- Session information

Depending on the software version running on your switch and the type of switch you have, additional account information may be displayed.

You must have administrator privileges to delete a user account. There must be one administrator account on the system; the command will fail if an attempt is made to delete the last administrator account on the system.

Example

The following command deletes account John2:

```
delete account john2
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

disable clipaging

```
disable clipaging
```

Description

Disables pausing at the end of each show screen.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

The command line interface (CLI) is designed for use in a VT100 environment. Most show command output will pause when the display reaches the end of a page. This command disables the pause mechanism and allows the display to print continuously to the screen.

NOTE

Press [q] and then press [Return] to force a pause when CLI paging is disabled.

To view the status of CLI paging on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for CLI paging.

Example

The follow command disables clipaging and allows you to print continuously to the screen:

```
disable clipaging
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on all platforms.

disable idletimeouts

```
disable idletimeouts
```

Description

Disables the timer that disconnects idle sessions from the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

When idle time-outs are disabled, console sessions remain open until the switch is rebooted or you logoff. Telnet sessions remain open until you close the Telnet client.

To view the status of idle time-outs on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for idle time-outs.

Example

The following command disables the timer that disconnects all sessions to the switch:

```
disable idletimeouts
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

enable clipaging

```
enable clipaging
```

Description

Enables the pause mechanism and does not allow the display to print continuously to the screen.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

The command line interface (CLI) is designed for use in a VT100 environment. Most show command output will pause when the display reaches the end of a page.

To view the status of CLI paging on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for CLI paging.

If CLI paging is enabled and you use the `show tech-support` command to diagnose system technical problems, the CLI paging feature is disabled.

Example

The following command enables clipaging and does not allow the display to print continuously to the screen:

```
enable clipaging
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on all platforms.

enable idletimeouts

```
enable idletimeouts
```

Description

Enables a timer that disconnects Telnet and console sessions after 20 minutes of inactivity.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

You can use this command to ensure that a Telnet, HTTP or console session is disconnected if it has been idle for the required length of time. This ensures that there are no hanging connections.

To view the status of idle time-outs on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for idle time-outs.

In ExtremeWare version 6.2 or later, you can configure the length of the time-out interval.

Example

The following command enables a timer that disconnects any Telnet, HTTP, and console sessions after 20 minutes of inactivity:

```
enable idletimeouts
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

enable license

```
enable license [basic_L3 | advanced_L3 | full_L3 ] <license_key>
```

Description

Enables a particular software feature license.

Syntax Description

basic_L3	Specifies a basic L3 license. (4.x only)
advanced_L3	Specifies an advanced L3 license. (4.x only)
full_L3	Specifies a full L3 license. (6.0, 6.1)
license_key	Specifies your software license key.

Default

N / A

Usage Guidelines

Specify `license_key` as an integer.

The `unconfig switch all` command does not clear licensing information. This feature cannot be disabled after the license has been enabled on the switch.

Depending on the software version running on your switch, and the type of switch you have, only the license parameters applicable to your software or switch can be used.

To view the type of license you are currently running on the switch, use the `show switch` command. The license key number is not displayed, but the type of license is displayed in the `show switch` output. The type of license is displayed after the system name, system location, system contact, and system MAC address.

Example

The following command enables a full L3 license on the switch:

```
enable license fullL3
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

history

```
history
```

Description

Displays a list of the previous 49 commands entered on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

ExtremeWare “remembers” the last 49 commands you entered on the switch. Use the `history` command to display a list of these commands.

Example

The following command displays the previous 49 commands entered on the switch:

```
history
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

nslookup

```
nslookup <hostname>
```

Description

Displays the IP address of the requested host.

Syntax Description

hostname	Specifies a hostname.

Default

N/A.

Usage Guidelines

None.

Example

The following command looks up the IP address of a computer with the name of *bigserver.xyz_inc.com*:

```
nslookup bigserver.xyz_inc.com
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

ping

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

Description

Enables you to send User Datagram Protocol (UDP) or Internet Control Message Protocol (ICMP) echo messages or to a remote IP device.

Syntax Description

udp	Specifies that the ping request should use UDP instead of ICMP (6.2 or higher).
continuous	Specifies that UDP or ICMP echo messages to be sent continuously. This option can be interrupted by pressing any key.
start_size	Specifies the size, in bytes, of the packet to be sent, or (in 6.1 or higher) the starting size if incremental packets are to be sent.
end_size	Specifies the maximum size, in bytes, of the packet to be sent in the UDP or ICMP request. When both the start_size and end_size are specified, ICMP requests are transmitted using 1 byte increments, per packet. (6.1 and higher)
ipaddress	Specifies the IP address of the host.
hostname	Specifies the name of the host. (6.1 and higher)
src_ipaddress	Uses the specified source address in the ICMP packet. If not specified, the address of the transmitting interface is used. (6.1 and higher)
record-route	Decodes the list of recorded routes and displays them when the ICMP echo reply is received. (6.1 and higher)

Default

N / A.

Usage Guidelines

The `ping` command is used to test for connectivity to a specific host.

The `ping` command is available for both the user and administrator privilege level.

If a `ping` request fails, the switch continues to send `ping` messages until interrupted. Press any key to interrupt a `ping` request.

For version 6.1:

- You must configure DNS in order to use the `hostname` option.

For version 6.2:

- If you specify UDP as the protocol, the `from <source>` and `with <record-route>` options are not supported.

Example

The following command enables continuous ICMP echo messages to be sent to a remote host:

```
ping continuous 123.45.67.8
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 6.1 to support the `hostname`, `from`, and `with record-route` parameters, and incremental packets.

This command was modified in ExtremeWare 6.2 to support UDP.

Platform Availability

This command is available on all platforms.

reboot

```
reboot {time <date> <time> | cancel} {slot <slot>}
```

Description

Reboots the switch at a specified date and time.

For PoS and MPLS modules:

Reboots the switch or the PoS or MPLS module in the specified slot at a specified date and time.

Syntax Description

date	Specifies a reboot date in `mm/dd/yyyy` format.
time	Specifies a reboot time in `hh:mm:ss` format.
cancel	Cancels a previously scheduled reboot.
slot	Specifies the slot where the PoS or MPLS module is installed. (6.2)

Default

N / A.

Usage Guidelines

If you do not specify a reboot time, the switch will reboot immediately following the command, and any previously scheduled reboots are cancelled. To cancel a previously scheduled reboot, use the `cancel` option.

The slot <slot> option is added to the command to make it possible to reboot a module in a specific slot. When you specify this option, the command applies to the PoS or MPLS module in the specified slot, rather than to the switch.

Example

The following command reboots the switch at 8:00 AM on April 15, 2002:

```
reboot 04/15/2002 08:00:00
```

The following command reboots the MPLS module in slot number 5:

```
reboot time 10/04/2001 10,46,00 slot 5
```

History

This command was first available in ExtremeWare 2.0.

This command is also available in ExtremeWare IP Technology Services Releases based on ExtremeWare 6.1.8 to support MPLS, ARM, and PoS modules.

Platform Availability

This command is available on all platforms.

run msm-failover

```
run msm-failover
```

Description

Causes a user-specified MSM failover.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command causes a user-specified MSM failover:

```
run msm-failover
```

History

This command was first available in ExtremeWare 6.2.2.

Platform Availability

This command is available on the BlackDiamond only.

show accounts

```
show accounts
```

Description

Displays user account information for all users on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

You need to create a user account using the `create account` command before you can display user account information.

To view the accounts that have been created, you must have administrator privileges.

The `show accounts` command displays the following information in a tabular format:

- User Name—The name of the user. This list displays all of the users who have access to the switch.
- Access—The SNMP community strings. This may be listed as R/W for read/write or RO for read only.
- Login OK—The number of logins that are okay.
- Failed—The number of failed logins.

Depending on the software version running on your switch, additional or different account information may be displayed.

Example

The following command displays user account information on the switch:

```
show accounts
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

show banner

```
show banner
```

Description

Displays the user-configured banner string.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to view the banner that is displayed before the login prompt.

Example

The following command displays the switch banner:

```
show banner
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

show dns-client

```
show dns-client
```

Description

Displays the DNS configuration.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the DNS configuration:

```
show dns-client
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

show switch

```
show switch
```

Description

Displays the current switch information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

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. This way, statistics can help you get the best out of your network.

The `show switch` command displays:

- sysName, sysLocation, sysContact
- MAC address
- System mode
- RED configuration
- DLCS state
- Recovery mode
- Watchdog state
- Current date, time, system uptime, and time zone configuration
- Operating environment (temperature, fans, and power supply status)
- Software image information (primary/secondary image, date/time, version)
- NVRAM configuration information (primary/secondary configuration, date/time, size, version)
- Any scheduled reboot information
- Scheduled download information
- PACE configuration information
- Software licensing information
- MSM information (BlackDiamond only)

This information may be useful for your technical support representative if you have a problem.

Depending on the software version running on your switch, additional or different switch information may be displayed.

Example

The following command displays current switch information:

```
show switch
```

Output from this command will look similar to the following:

```
SysName:          Summit1iTx
SysLocation:
SysContact:       support@extremenetworks.com, +1 888 257 3000
System MAC:       00:01:30:70:06:00

License:          Full L3 + Security
System Mode:      802.1Q EtherType is 8100 (Hex).   CPU Tx-Priority = High

RED Probability:  0
DLCS:             Disabled

Recovery Mode:    None
System Watchdog:  Enabled

Current Time:     Mon Feb 25 01:15:30 2002
Timezone:         [Auto DST Enabled] GMT Offset: 0 minutes, name is not set.
                  DST of 60 minutes is currently not in effect, name is not set.
                  DST begins every first Sunday April at 2:00
                  DST ends every last Sunday October at 2:00
Boot Time:        Mon Feb 25 00:37:14 2002
Next Reboot:      None scheduled
Timed Upload:     None scheduled
Timed Download:   None scheduled

Temperature:      Normal.    All fans are operational.
Power supply:     PSU-A OK,   PSU-B powered off

Image Selected:   Primary
Image Booted:     Primary

Primary EW Ver:   6.2.1b17
Secondary EW Ver: 6.1.5b20

Config Selected:  Primary
Config Booted:    Primary

Primary Config:   Created by EW Version: 6.2.1 Build 17 [32]
                  3472 bytes saved on Fri Feb 15 02:05:05 2002
Secondary Config: Created by EW Version: 6.1.5 Build 20 [32]
                  3260 bytes saved on Mon Dec 17 14:11:52 2001
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

traceroute

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

Description

Enables you to trace the routed path between the switch and a destination endstation.

Syntax Description

ipaddress	Specifies the IP address of the destination endstation.
hostname	Specifies the hostname of the destination endstation. (4.x, 6.0 and higher)
from	Uses the specified source address in the ICMP packet. If not specified, the address of the transmitting interface is used. (6.1 and higher)
ttl	Configures the switch to trace up to the time-to-live number of the switch. (6.1 and higher)
port	Specifies the UDP port number. (6.1 and higher)

Default

N/A.

Usage Guidelines

To use the `hostname` parameter, you must first configure DNS.

Each router along the path is displayed.

Example

The following command enables the traceroute function to a destination of 123.45.67.8:

```
traceroute 123.45.67.8
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 6.1 to support the `from`, `ttl`, and `port` parameters.

This command was modified in ExtremeWare 4.0 to support the `hostname` parameter.

Platform Availability

This command is available on all platforms.

3 Commands for Managing the Switch

This chapter describes:

- Commands for configuring Simple Network Management Protocol (SNMP) parameters on the switch
- Commands for managing the switch using SNMP, Telnet, SSH2, and web access
- Commands related to switch user authentication through a RADIUS client
- Commands related to switch user authentication through TACACS+

SNMP

Any network manager running the Simple Network Management Protocol (SNMP) can manage the switch, if 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 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. Entries in this list can be created, modified, and deleted using the RMON2 trapDestTable MIB variable, as described in RFC 2021.
- Authorized managers — An authorized manager can be either a single network management station, or a range of addresses (for example, a complete subnet) specified by a prefix and a mask. The switch can have a maximum of eight authorized managers.
- Community strings — The community strings allow a simple method of authentication between the switch and the remote Network Manager. The default read-only community string is *public*. The default read-write community string is *private*. The community strings for all authorized trap receivers must be configured on the switch for the trap receiver to receive switch-generated traps.
- 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).
- System location (optional) — Using the system location field, you can enter an optional location for this switch.

The following can also be configured on the switch for version 6.0 and higher:

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

Telnet

Telnet allows you to access the switch remotely using TCP/IP through one of the switch ports or a workstation with a Telnet facility. If you access the switch via Telnet, you will use the command line interface (CLI) to manage the switch and modify switch configurations.

SSH

Secure Shell 2 (SSH2) is a feature of ExtremeWare that allows you to encrypt Telnet session data between a network administrator using SSH2 client software and the switch, or to send encrypted data from the switch to an SSH2 client on a remote system. Image and configuration files may also be transferred to the switch using the Secure Copy Protocol 2 (SCP2)

ExtremeWare Vista

ExtremeWare Vista is a device management software running in the switch that allows you to access the switch over a TCP/IP network using a standard web browser. ExtremeWare Vista provides a subset of the CLI commands available for configuring and monitoring the switch. If a particular command is not available using ExtremeWare Vista, you must use the CLI to access the desired functionality.

User Authentication

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

Extreme switches are also capable of sending RADIUS accounting information. You can configure RADIUS accounting servers to be the same as the authentication servers, but this is not required.

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.

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.

Network Login

Network login has two modes of operation:

- Campus mode, 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. Campus mode requires a DHCP server and a RADIUS server configured for Extreme network login.
- ISP mode, used when the port and VLAN used will remain constant. All network settings are configured for that VLAN.

Denial of Service

You can configure ExtremeWare to protect your network in the event of a denial of service attack. During a typical denial of service attack, the CPU on the switch gets flooded with packets from multiple attackers, potentially causing the switch to fail. To protect against this type of attack, you can configure the software so that when the number of packets received is more than the configured threshold limit of packets per second, a hardware ACL is enabled.

config cpu-dos-protect

```
config cpu-dos-protect [alert-threshold <packets per second>]
[notice-threshold <packets per second>] [timeout <seconds>] [messages [on |
off]] [filter-precedence <number>]
```

Description

Configures denial of service protection.

Syntax Description

alert-threshold	Configures the number of packets per second that the switch needs to recieve on a port for an ACL to be enabled. Range is 150 to 100,000 packets per second. Default is 4000.
notice-threshold	Configures the number of packets per second that the switch needs to receive on a port for messages to be logged. Range is 150 to 100,000 packets per second. Default is 4000.
timeout	Configures a duration in seconds. Range is 2 to 300 seconds. Default is 15.
messages	Configures messaging to be on or off. Default is on.
filter-precedence	Configures the access list precedence. Default is 10.

Default

The option defaults are:

- alert-threshold—4000
- notice-threshold—4000.
- timeout—15
- messages—on
- filter-precedence—10

Usage Guidelines

None.

Example

The following command configures denial of service protection to be invoked when 3000 or more packets per second are received by a port on the switch. This command configures logging to occur when the number of packets per second that the switch receives is 2000, the timeout is 15 seconds, and messages are on:

```
config cpu-dos-protect alert-threshold 3000 notice-threshold 2000 timeout 15 messages
on filter-precedence 10
```

History

This command was first available in ExtremeWare 6.2.2

Platform Availability

This command is available on all platforms.

config radius server

```
config radius [primary | secondary] server [<ipaddress> | <hostname>]
{<udp_port>} client-ip [<ipaddress>]
```

Description

Configures the primary and secondary RADIUS authentication server.

Syntax Description

primary	Configures the primary RADIUS authentication server.
secondary	Configures the secondary RADIUS authentication server.
ipaddress	The IP address of the server being configured.
hostname	The host name of the server being configured.
udp_port	The UDP port to use to contact the RADIUS authentication server.
ipaddress	The IP address used by the switch to identify itself when communicating with the RADIUS authentication server.

Default

The default UDP port setting is 1645.

Usage Guidelines

Use this command to specify RADIUS server information.

Use of the <hostname> parameter requires that DNS be enabled.

The RADIUS server defined by this command is used for user name authentication and CLI command authentication.

Example

The following command configures the primary RADIUS server on host `radius1` using the default UDP port (1645) for use by the RADIUS client on switch 10.10.20.30:

```
config radius primary server radius1 client-ip 10.10.20.30
```

History

This command was first available in ExtremeWare 4.1.

Platform Availability

This command is available on all platforms.

config radius shared-secret

```
config radius [primary | secondary] shared-secret {encrypted} [<string>]
```

Description

Configures the authentication string used to communicate with the RADIUS authentication server.

Syntax Description

primary	Configures the authentication string for the primary RADIUS server.
secondary	Configures the authentication string for the secondary RADIUS server.
encrypted	Indicates that the secret should be encrypted
string	The string to be used for authentication.

Default

Unconfigured.

Usage Guidelines

The secret must be the same between the client switch and the RADIUS server.

The RADIUS server must first be configured for use with the switch as a RADIUS client.

Example

The following command configures the shared secret as "purplegreen" on the primary RADIUS server:

```
config radius primary shared-secret purplegreen
```

History

This command was first available in ExtremeWare 4.1.

Platform Availability

This command is available on all platforms.

config radius-accounting server

```
config radius-accounting [primary | secondary] server [<ipaddress> |
<hostname>] {<udp_port>} client-ip [<ipaddress>]
```

Description

Configures the RADIUS accounting server.

Syntax Description

primary	Configure the primary RADIUS accounting server.
secondary	Configure the secondary RADIUS accounting server.
ipaddress	The IP address of the accounting server being configured.
hostname	The host name of the accounting server being configured.
udp_port	The UDP port to use to contact the RADIUS accounting server.
ipaddress	The IP address used by the switch to identify itself when communicating with the RADIUS accounting server.

Default

The default UDP port setting is 1646.

Usage Guidelines

Use this command to specify the radius accounting server.

The accounting server and the RADIUS authentication server can be the same.

Use of the <hostname> parameter requires that DNS be enabled.

Example

The following command configures RADIUS accounting on host `radius1` using the default UDP port (1646) for use by the RADIUS client on switch 10.10.20.30:

```
config radius-accounting primary server radius1 client-ip 10.10.20.30
```

History

This command was first available in ExtremeWare 4.1.

Platform Availability

This command is available on all platforms.

config radius-accounting shared-secret

```
config radius-accounting [primary | secondary] shared-secret {encrypted}
[<string>]
```

Description

Configures the authentication string used to communicate with the RADIUS accounting server.

Syntax Description

primary	Configures the authentication string for the primary RADIUS accounting server.
secondary	Configures the authentication string for the secondary RADIUS accounting server.
encrypted	Indicates that the secret should be encrypted
string	The string to be used for authentication.

Default

Unconfigured.

Usage Guidelines

The secret must be the same between the client switch and the RADIUS accounting server.

Example

The following command configures the shared secret as "purpleaccount" on the primary RADIUS accounting server:

```
config radius primary shared-secret purpleaccount
```

History

This command was first available in ExtremeWare 4.1.

Platform Availability

This command is available on all platforms.

config snmp access-profile readonly

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

Description

Assigns an access profile that limits which stations have read-only access to the switch.

Syntax Description

access-profile	Specifies a user defined access profile.
none	Cancels a previously configured access profile.

Default

All users have access until an access profile is created and specified.

Usage Guidelines

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

You must create and configure an access profile before you can use this command. You create an access profile using the `create access-profile` command. You configure an access profile using the `config access-profile` command.

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

Read community strings provide read-only access to the switch. The default read-only community string is public. A total of eight community strings (read-only and read/write) 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.

To view the SNMP read-only access communities configured on the switch, use the `show management` command. The `show management` command displays information about the switch including the names and the number of read-only communities configured on the switch.

To restore defaults to all SNMP-related entries, including the SNMP parameters modified using the `config snmp access-profile readonly` command, use the `unconfig management` command.

Example

The following command allows the user defined access profile *admin* read-only access to the switch:

```
config snmp access-profile readonly admin
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

config snmp access-profile readwrite

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

Description

Assigns an access profile that limits which stations have read/write access to the switch.

Syntax Description

access-profile	Specifies a user defined access profile.
none	Cancels a previously configured access profile.

Default

All users have access until an access profile is specified.

Usage Guidelines

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

You must create and configure an access profile before you can use this command. You create an access profile using the `create access-profile` command. You configure an access profile using the `config access-profile` command.

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

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 (read-only and read/write) 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.

To view the SNMP read/write access communities configured on the switch, use the `show management` command. The `show management` command displays information about the switch including the names and the number of read/write communities configured on the switch.

To restore defaults to all SNMP-related entries, including the SNMP parameters modified using the `config snmp access-profile readwrite` command, use the `unconfig management` command.

Example

The following command allows the user defined access profile *management* read/write access to the switch:

```
config snmp access-profile readwrite management
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

config snmp add

```
config snmp add <ipaddress> {<mask>}
```

Description

Adds the IP address or a set of IP addresses of an SNMP management station to the access list.

Syntax Description

ipaddress	Specifies an IP address to be added to the access list.
mask	Specifies a subnet mask.

Default

All users are allowed access until an IP address or an IP address group is specified.

Usage Guidelines

Support for this command was discontinued in ExtremeWare 6.x.

Do not configure access profiles if you use this command. This command sets access parameters for a specific IP address or an IP address group. If you configure access profiles, you may disrupt the settings specified with this command.

After you add an IP address or an IP address group, you restrict access to that management station. Only those associated with the added station are granted access.

A maximum of 32 entries, which include individual IP addresses or IP address groups, can be specified.

An authorized manager can be either a single network management station, or a range of addresses (for example, a complete subnet) specified by a prefix and a mask. The switch can have a maximum of eight authorized managers.

To restore defaults to all SNMP-related entries, including the SNMP parameters modified using the `config snmp add ipaddress` command, use the `unconfig management` command.

Example

The following command adds an SNMP management station with an IP address of 123.45.67.8 to the access list:

```
config snmp add 123.45.67.8
```

After it has been configured, access is restricted to the specified SNMP management station(s) only.

History

This command was first available in ExtremeWare 2.0.

Support for this command was discontinued in ExtremeWare 6.0.

Platform Availability

This command is available only on platforms based on the Summit chipset.

config snmp add community

```
config snmp add community [readonly | readwrite] {encrypted} <alphanumeric
string>
```

Description

Adds an SNMP read or read/write community string.

Syntax Description

readonly	Specifies read-only access to the system.
readwrite	Specifies read and write access to the system.
encrypted	Specifies encryption, for use only by the switch when uploading or downloading a configuration. Should not be used through the CLI.
alphanumeric string	Specifies an SNMP community string name. See “Usage Guidelines” for more information.

Default

The default read-only community string is *public*. The default read/write community string is *private*.

Usage Guidelines

Community strings provide a simple method of authentication between a switch and a remote network manager. 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*.

An authorized trap receiver must be configured to use the correct community strings on the switch for the trap receiver to receive switch-generated traps. In some cases, it may be useful to allow multiple community strings so that all switches and trap receivers are not forced to use identical community strings. The `config snmp add community` command allows you to add multiple community strings in addition to the default community string.

A total of 15 community strings can be configured on the switch. An SNMP community string can contain up to 127 characters.

To change the value of the default read/write and read-only community strings, use the `config snmp community` command.

The `encrypted` option is intended for use by the switch when generatinge an ASCII configuration file (using the `upload configuration` command), or parsing a switch-generated configuration (using the `download configuration` command). Do not select the encrypted option in the CLI.

Example

The following command adds a read/write community string with the value *extreme*:

```
config snmp add community readwrite extreme
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all "*i*"-series platforms.

config snmp add trapreceiver

```
config snmp add trapreceiver <ip address> {port <udp_port>} community
<community string> {from <source ip address>} {mode [enhanced | standard]}
```

Description

Adds the IP address of a specified trap receiver to the trap receiver list.

Syntax Description

ip address	Specifies an SNMP trap receiver IP address.
udp_port	Specifies a UDP port to which the trap should be sent. Default is 162.
community string	Specifies the community string of the trap receiver.
source ip address	Specifies the IP address of a VLAN to be used as the source address for the trap
enhanced	Specifies enhanced traps, which contain extra varbinds at the end.
standard	Specifies standard traps, which do not contrain the extra varbinds.

Default

Trapreceivers are in enhanced mode by default.

Usage Guidelines

The IP address can be unicast, multicast, or broadcast.

An authorized trap receiver can be one or more network management stations on your network. Authorized trap receivers must be configured on the switch for the trap receiver to receive switch-generated traps. The switch sends SNMP traps to all trap receivers. Entries in this list can be created, modified, and deleted using the RMON2 trapDestTable MIB variable, as described in RFC 2021.

To view the SNMP trap receivers configured on the switch, use the `show management` command. The `show management` command displays information about the switch including the destination and community of the SNMP trap receivers configured on the switch.

To restore defaults to all SNMP-related entries, including the SNMP parameters modified using the `config snmp add trapreceiver` command, use the `unconfig management` command.

For version 6.0 and higher:

- A maximum of sixteen trap receivers can be configured for each switch.

For version 4.x:

- A maximum of six trap receivers can be configured for each switch.

Example

The following command adds the IP address 10.101.0.100 as a trap receiver with community string *purple*:

```
config snmp add trapreceiver 10.101.0.100 community purple
```

The following command adds the IP address 10.101.0.105 as a trap receiver with community string *green*, using port 3003:

```
config snmp add trapreceiver 10.101.0.105 port 3003 community green
```

The following command adds the IP address 10.101.0.105 as a trap receiver with community string *blue*, and IP address 10.101.0.25 as the source:

```
config snmp add trapreceiver 10.101.0.105 community blue from 10.101.0.25
```

History

This command was first available in ExtremeWare 1.0.

This command was modified in ExtremeWare 6.2.1 to support the `port`, `community`, and source (`from`) options.

This command was modified in ExtremeWare 6.2.2 to add the `mode` options.

Platform Availability

This command is available on all platforms.

config snmp community

```
config snmp community [readonly | readwrite] {encrypted} <alphanumeric
string>
```

Description

Configures the value of the default SNMP read or read/write community string.

Syntax Description

readonly	Specifies read-only access to the system.
readwrite	Specifies read and write access to the system.
encrypted	Specifies encryption, for use only by the switch when uploading or downloading a configuration. Should not be used through the CLI.
alphanumeric string	Specifies an SNMP community string name. See “Usage Guidelines” for more information.

Default

The default read-only community string is *public*. The default read/write community string is *private*.

Usage Guidelines

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*.

It is recommended that you change the values of the default read/write and read-only community strings. You use the `config snmp community` command to change the value of the default community strings. An SNMP community string can contain up to 127 characters.

The `encrypted` option is inteded for use by the switch when generatinge an ASCII configuration file (using the `upload configuration` command), or parsing a switch-generated configuration (using the `download configuration` command). Do not select the encrypted option in the CLI.

For version 6.2:

- A total of fifteen community strings can be configured on the switch. You can add additional community strings (in addition to the default community stings) using the `config snmp add community` command.

Example

The following command sets the read/write community string to *extreme*:

```
config snmp community readwrite extreme
```

History

This command was first available in ExtremeWare 1.0.

Platform Availability

This command is available on all platforms.

config snmp delete

```
config snmp delete [<ip address> {<mask>} | all]
```

Description

Deletes an IP address or range of IP addresses of a specified SNMP management station or all SNMP management stations.

Syntax Description

ip address	Specifies an SNMP management IP address.
mask	Specifies an optional subnet mask.
all	Specifies all SNMP management IP addresses.

Default

N/A.

Usage Guidelines

Support for this command was discontinued in ExtremeWare 6.x.

If you delete all addresses, any machine can have SNMP access to the switch.

Example

The following command deletes IP address 123.45.67.8 from an SNMP station:

```
config snmp delete 123.45.67.8
```

History

This command was first available in ExtremeWare 2.0.

Support for this command was discontinued in ExtremeWare 6.0.

Platform Availability

This command is available only on platforms based on the Summit chipset.

config snmp delete community

```
config snmp delete community [readonly | readwrite] {encrypted} [all |
<alphanumeric string>]
```

Description

Deletes an SNMP read or read/write community string.

Syntax Description

readonly	Specifies read-only access to the system.
readwrite	Specifies read and write access to the system.
encrypted	Specifies an encrypted option.
all	Specifies all of the SNMP community stings.
alphanumeric string	Specifies an SNMP community string name. See “Usage Guidelines” for more information.

Default

The default read-only community string is *public*. The default read/write community string is *private*.

Usage Guidelines

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 15 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.

It is recommended that you change the defaults of the read/write and read-only community strings.

Use the `config snmp add` command to configure an authorized SNMP management station.

The `encrypted` option should only be used by the switch to generate an ASCII configuration (using the `upload configuration` command), and parsing a switch-generated configuration (using the `download configuration` command). Do not select the encrypted option in the CLI.

For version 6.0 and 6.1:

- A total of eight community strings can be configured on the switch.

For version 4.x:

- SNMP community strings can contain up to 126 characters.

For version 2.0:

- The `add` parameter is included in the command syntax. It is available only in version 2.0.
- SNMP community strings can contain up to 127 characters.

Example

The following command adds a read/write community string named *extreme*:

```
config snmp add community readwrite extreme
```

History

This command was first available in ExtremeWare 2.0.

Support for the add parameter was discontinued in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

config snmp delete trapreceiver

```
config snmp delete trapreceiver [<ip address>
{community <community string>} | all]
```

Description

Deletes a specified trap receiver or all authorized trap receivers.

Syntax Description

ip address	Specifies an SNMP trap receiver IP address.
community string	Specifies an SNMP community string.
all	Specifies all SNMP trap receiver IP addresses.

Default

N / A.

Usage Guidelines

Use this command to delete a trap receiver of the specified IP address, or all authorized trap receivers.

If a trap receiver has been added multiple times with different community strings, the `community` option specifies that only the trap receiver entry with the specified community string should be removed.

Example

The following command deletes the trap receiver 10.101.0.100 from the trap receiver list:

```
config snmp delete trapreceiver 10.101.0.100
```

The following command deletes entries in the trap receiver list for 10.101.0.100 with community string public:

```
config snmp delete trapreceiver 10.101.0.100 community public
```

Any entries for this IP address with a different community string will not be affected.

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 6.2.1 to support the `community` option.

Platform Availability

This command is available on all platforms.

config snmp sysContact

```
config snmp syscontact <alphanumeric string>
```

Description

Configures the name of the system contact.

Syntax Description

alphanumeric string	Specifies a system contact name.

Default

N/A.

Usage Guidelines

The system contact is a text field that enables you to enter the name of the person(s) responsible for managing the switch. A maximum of 255 characters is allowed.

To view the name of the system contact listed on the switch, use the `show switch` command. The `show switch` command displays switch statistics including the name of the system contact.

To restore defaults to all SNMP-related entries, including the SNMP parameters modified using the `config snmp syscontact <alphanumeric string>` command, use the `unconfig management` command.

Example

The following command defines FredJ as the system contact:

```
config snmp syscontact fredj
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config snmp sysLocation

```
config snmp syslocation <alphanumeric string>
```

Description

Configures the location of the switch.

Syntax Description

alphanumeric string	Specifies the switch location.

Default

N / A.

Usage Guidelines

Use this command to indicate the location of the switch. A maximum of 255 characters is allowed.

To view the location of the switch on the switch, use the `show switch` command. The `show switch` command displays switch statistics including the location of the switch.

To restore defaults to all SNMP-related entries, including the SNMP parameters modified using the `config snmp syslocation <alphanumeric string>` command, use the `unconfig management` command.

Example

The following command configures a switch location name on the system:

```
config snmp syslocation englab
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config snmp sysName

```
config snmp sysname <alphanumeric string>
```

Description

Configures the name of the switch.

Syntax Description

alphanumeric string	Specifies a device name.

Default

The default sysname is the model name of the device (for example, Summit1).

Usage Guidelines

You can use this command to change the name of the switch. A maximum of 32 characters is allowed. The sysname appears in the switch prompt.

To view the name of the system listed on the switch, use the show switch command. The show switch command displays switch statistics including the name of the system.

To restore defaults to all SNMP-related entries, including the SNMP parameters modified using the config snmp sysname <alphanumeric string> command, use the unconfig management command.

Example

The following command names the switch:

```
config snmp sysname engineeringlab
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config sntp-client server

```
config sntp-client [primary | secondary] server [<ipaddress> | <host_name>]
```

Description

Configures an NTP server for the switch to obtain time information.

Syntax Description

primary	Specifies a primary server name.
secondary	Specifies a secondary server name.
ipaddress	Specifies an IP address.
host_name	Specifies a host name.

Default

N / A.

Usage Guidelines

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. 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.

Example

The following command configures a primary NTP server:

```
config sntp-client primary server 10.1.2.2
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

config sntp-client update-interval

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

Description

Configures the interval between polls for time information from SNTP servers.

Syntax Description

seconds	Specifies an interval in seconds.

Default

64 seconds.

Usage Guidelines

None.

Example

The following command configures the interval timer:

```
config sntp-client update-interval 30
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

config ssh2

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

Description

Generates the Secure Shell 2 (SSH2) host key.

Syntax Description

pregenerated	Indicates that the SSH2 authentication key has already been generated. The user will be prompted to enter the existing key.

Default

The switch generates a key for each SSH2 session.

Usage Guidelines

Secure Shell 2 (SSH2) is a feature of ExtremeWare that allows you to encrypt Telnet session data between a network administrator using SSH2 client software and the switch, or to send encrypted data from the switch to an SSH2 client on a remote system. Image and configuration files may also be transferred to the switch using the Secure Copy Protocol (SCP) or the Secure File Transfer Protocol (SFTP).

Before you can enable SSH2, you must first obtain a security license from Extreme Networks. After you receive the license, you must enable SSH2 and generate a host key. To enable SSH2, use the `enable ssh2` command. To generate an SSH2 host key, use the `config ssh2 key` command.

An authentication key must be generated before the switch can accecpt incoming SSH2 sessions. This can be done automatically by the switch, or you can enter a previously generated key.

If you elect to have the key generated, you are prompted to enter a set of random characters to be used in generating the key. The key generation process takes approximately ten minutes, and cannot be canceled after it has started. Once the key has been generated, you should save your configuration to preserve the key.

To use a key that has been previously created, use the `pregenerated` keyword. You are prompted to enter the pregenerated key. You can use the `show config` command to list the previously generated key, and then copy and paste it after the prompt from the `config ssh2 key pregenerated` command.

The key generation process generates the SSH2 private host key. The SSH2 public host key is derived from the private host key, and is automatically transmitted to the SSH2 client at the beginning of an SSH2 session.

Example

The following command generates an authentication key for the SSH2 session:

```
config ssh2 key
```

The command responds with the following messages:

```
WARNING: Generating new server host key
This will take approximately 10 minutes and cannot be canceled.
Continue? (y/n)
```

If you respond yes, the command prompts as follows:

```
Enter some random characters.  End with a newline
```

Type in a series of random characters, and then press the Enter or Return key. The key generation process will then proceed.

To configure an SSH2 session using a previously generated key, use the following command:

```
config ssh2 key pregenerated
```

The command responds with the following message:

```
Please enter the server key
```

Enter the previously-generated key (you can copy and paste it from the saved configuration file).

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on all platforms.

config tacacs server

```
config tacacs [primary | secondary] server [<ipaddress> | <hostname>]
{<tcp_port>} client-ip <ipaddress>
```

Description

Configures the server information for a TACACS+ authentication server.

Syntax Description

primary	Configures the primary TACACS+ server.
secondary	Configures the secondary TACACS+ server.
ipaddress	The IP address of the TACACS+ server being configured.
hostname	The host name of the TACACS+ server being configured.
tcp_port	The TCP port to use to contact the TACACS+ server.
ipaddress	The IP address used by the switch to identify itself when communicating with the TACACS+ server.

Default

TACACS+ uses TCP port 49.

Usage Guidelines

Configure the server information for a TACACS+ server.

To remove a server, use the address 0.0.0.0 as the IP address of the server to be configured.

Use of the <hostname> parameter requires that DNS be enabled.

Example

The following command configures server tacacs1 as the primary TACACS+ server for client switch 10.10.20.35:

```
config tacacs primary server tacacs1 client-ip 10.10.20.35
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the "*i*" series platforms.

config tacacs shared-secret

```
config tacacs [primary | secondary] shared-secret {encrypted} <string>
```

Description

Configures the shared secret string used to communicate with the TACACS+ authentication server.

Syntax Description

primary	Configures the authentication string for the primary TACACS+ server.
secondary	Configures the authentication string for the secondary TACACS+ server.
encrypted	Indicates that the secret should be encrypted.
string	The string to be used for authentication.

Default

N/A.

Usage Guidelines

The secret must be the same between the client switch and the TACACS+ server.

Example

The following command configures the shared secret as "purplegreen" on the primary TACACS+ server:

```
config tacacs-accounting primary shared-secret purplegreen
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the "*i*" series platforms.

config tacacs-accounting server

```
config tacacs-accounting [primary | secondary] server [<ipaddress> |
<hostname>] {<udp_port>} client-ip <ipaddress>
```

Description

Configures the TACACS+ accounting server.

Syntax Description

primary	Configures the primary TACACS+ accounting server.
secondary	Configures the secondary TACACS+ accounting server.
ipaddress	The IP address of the TACACS+ accounting server being configured.
hostname	The host name of the TACACS+ accounting server being configured.
tcp_port	The TCP port to use to contact the TACACS+ server.
ipaddress	The IP address used by the switch to identify itself when communicating with the TACACS+ accounting server.

Default

Unconfigured.

Usage Guidelines

You can use the same TACACS+ server for accounting and authentication.

To remove a server, use the address 0.0.0.0 as the IP address of the server to be configured.

Example

The following command configures server tacacs1 as the primary TACACS+ accounting server for client switch 10.10.20.35:

```
config tacacs-accounting primary server tacacs1 client-ip 10.10.20.35
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the "*i*" series platforms.

config tacacs-accounting shared-secret

```
config tacacs-accounting [primary | secondary] shared-secret {encrypted}
<string>
```

Description

Configures the shared secret string used to communicate with the TACACS+ accounting server.

Syntax Description

primary	Configures the authentication string for the primary TACACS+ accounting server.
secondary	Configures the authentication string for the secondary TACACS+ accounting server.
encrypted	Indicates that the secret should be encrypted.
string	The string to be used for authentication.

Default

N/A.

Usage Guidelines

Secret needs to be the same as on the TACACS+ server.

Example

The following command configures the shared secret as "tacacsaccount" on the primary TACACS+ accounting server:

```
config tacacs-accounting primary shared-secret tacacsaccount
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the *"i"* series platforms.

config vlan dhcp-address-range

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

Description

Configures a set of DHCP addresses for a VLAN.

Syntax Description

name	Specifies the VLAN on whose ports netlogin should be disabled.
ipaddress1	Specifies the first IP address in the DHCP address range to be assigned to this VLAN.
ipaddress2	Specifies the last IP address in the DHCP address range to be assigned to this VLAN.

Default

N / A.

Usage Guidelines

None.

Example

The following command allocates the IP addresses between 192.168.0.20 and 192.168.0.100 for use by the VLAN *temporary*:

```
config temporary dhcp-address-range 192.168.0.20 - 192.168.0.100
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan dhcp-lease-timer

```
config vlan <name> dhcp-lease-timer <lease-timer>
```

Description

Configures the timer value in seconds returned as part of the DHCP response.

Syntax Description

name	Specifies the VLAN on whose ports netlogin should be disabled.
lease-timer	Specifies the timer value, in seconds.

Default

N/A.

Usage Guidelines

The timer value is specified in seconds.

Example

The following command configures the DHCP lease timer value for VLAN *corp*:

```
config vlan corp dhcp-lease-timer <lease-timer>
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan dhcp-options

```
config vlan <name> dhcp-options [default-gateway | dns-server |
wins-server] <ipaddress>
```

Description

Configures the DHCP options returned as part of the DHCP response by a switch configured as a DHCP server.

Syntax Description

name	Specifies a VLAN name.
default-gateway	Specifies the router option.
dns-server	Specifies the Domain Name Server (DNS) option.
wins-server	Specifies the NetBIOS name server (NBNS) option.
ipaddress	The IP address associated with the specified option.

Default

N/A.

Usage Guidelines

None.

Example

The following command configures the DHCP server to return the IP address 10.10.20.8 as the router option:

```
config vlan <name> dhcp-options default-gateway 10.10.20.8
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan netlogin-lease-timer

```
config vlan <vlan_name> netlogin-lease-timer <lease-timer>
```

Description

Configures the timer value returned as part of the DHCP response for clients attached to network enabled ports.

Syntax Description

vlan_name	Specifies the VLAN to which this timer value applies.
lease-timer	Specifies the timer value, in seconds.

Default

30 seconds.

Usage Guidelines

The timer value is specified in seconds.

Example

The following command configures the timer value for VLAN *corp*:

```
config vlan corp netlogin-lease-timer <lease-timer>
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

disable cpu-dos-protect

```
disable cpu-dos-protect
```

Description

Disables denial of service protection.

Syntax Description

There are no arguments or variables for this command.

Default

Default is disabled.

Usage Guidelines

None.

Example

The following command disables denial of service protection.

```
disable cpu-dos-protect
```

History

This command was first available in ExtremeWare 6.2.2

Platform Availability

This command is available on all platforms.

disable dhcp ports vlan

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

Description

Disables DHCP on a specified port in a VLAN.

Syntax Description

portlist	Specifies the ports for which DHCP should be disabled.
vlan_name	Specifies the VLAN on whose ports DHCP should be disabled.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables DHCP for port 9 in VLAN *corp*:

```
disable dhcp ports 9 vlan corp
```

History

This command was first available in ExtremeWare 6.2.

disable netlogin ports

```
disable netlogin ports <portlist> vlan <vlan_name>
```

Description

Disables network login on a specified port in a VLAN.

Syntax Description

portlist	Specifies the ports for which netlogin should be disabled.
vlan_name	Specifies the VLAN on whose ports netlogin should be disabled.

Default

N/A.

Usage Guidelines

Network login must be disabled on a port before you can delete a VLAN that contains that port.

Example

The following command disables network login on port 9 in VLAN *corp*:

```
disable netlogin ports 9 vlan corp
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

disable radius

```
disable radius
```

Description

Disables the RADIUS client.

Syntax Description

This command has no arguments or variables.

Default

RADIUS authentication is disabled by default.

Usage Guidelines

None.

Example

The following command disables RADIUS authentication for the switch:

```
disable radius
```

History

This command was first available in ExtremeWare 4.1.

Platform Availability

This command is available on all platforms.

disable radius-accounting

```
disable radius-accounting
```

Description

Disables RADIUS accounting.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables RADIUS accounting for the switch:

```
disable radius-accounting
```

History

This command was first available in ExtremeWare 4.1.

Platform Availability

This command is available on all platforms.

disable snmp access

```
disable snmp access
```

Description

Disables SNMP on the switch.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Disabling SNMP access does not affect the SNMP configuration (for example, community strings). However, if you disable SNMP access, you will be unable to access the switch using SNMP.

Example

The following command disables SNMP access on the switch:

```
disable snmp access
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

disable snmp dot1dTpFdbTable

```
disable snmp dot1dpfdbtable
```

Description

Disables SNMP GetNext responses for the dot1dTpFdbTable in the BRIDGE-MIB.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

SNMP Get responses are not affected by this command.

To view the configuration of the dot1dTpFdp table on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state the dot1dTpFdp table.

To restore defaults to all SNMP-related entries, including the SNMP parameters modified using the `disable snmp dot1dtpfdbtable` command, use the `unconfig management` command.

Example

The following command disables the dot1dTPFdb table:

```
disable snmp dot1dtpfdbtable
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all platforms.

disable snmp traps

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

Description

Prevents SNMP traps from being sent from the switch.

Syntax Description

portlist	Specifies the ports for which traps should be disabled.
all	Specifies all ports.

Default

Enabled.

Usage Guidelines

This command does not clear the SNMP trap receivers that have been configured. The command prevents SNMP traps from being sent from the switch even if trap receivers are configured.

Example

The following command prevents SNMP traps from being sent from the switch to the trap receivers:

```
disable snmp traps
```

History

This command was first available in ExtremeWare 2.0.

This command was modified to include ports in ExtremeWare 6.2.2.

Platform Availability

This command is available on all platforms.

disable sntp-client

```
disable sntp-client
```

Description

Disables the SNTP client.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

SNTP can be used by the switch to update and synchronize its internal clock from a Network Time Protocol (NTP) server. After the SNTP client has been 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 (DST).

Example

The following command disables the SNTP client:

```
disable sntp-client
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

disable ssh2

```
disable ssh2
```

Description

Enables incoming SSH2 Telnet sessions.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

SSH2 session options (access profile and non-default port setting) are not saved when SSH2 is disabled.

To view the status of SSH2 Telnet sessions on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for SSH2 Telnet sessions.

Example

The following command disables the SSH2 feature:

```
disable ssh2
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on all "*i*" series platforms.

disable system-watchdog

```
disable system-watchdog
```

Description

Disables the system watchdog timer.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

The watchdog timer reboots the switch if the CPU becomes trapped in a processing loop. If the watchdog timer is executed, the switch captures information on the cause of the reboot and posts it to the system log.

Example

The following command disables the watchdog timer:

```
disable system-watchdog
```

History

This command was first available in ExtremeWare 6.1.9.

Platform Availability

This command is available on all "*i*" series platforms.

disable tacacs

```
disable tacacs
```

Description

Disables TACACS+ for authentication and authorization.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables TACACS+ authentication and authorization for the switch:

```
disable tacacs
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the "*i*" series platforms.

disable tacacs-accounting

```
disable tacacs-accounting
```

Description

Disables TACACS+ accounting.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables TACACS+ accounting:

```
disable tacacs-accounting
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the "*i*" series platforms.

disable tacacs-authorization

```
disable tacacs-authorization
```

Description

Disables CLI command authorization.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This disable CLI command authorization but leaves user authentication enabled.

Example

The following command disables TACACS+ CLI command authorization:

```
disable tacacs-authorization
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the "*i*" series platforms.

disable telnet

```
disable telnet
```

Description

Disables Telnet services on the system.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

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

Example

With administrator privilege, the following command disables Telnet services on the switch:

```
disable telnet
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

disable web

```
disable web
```

Description

Disables web access to the switch.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

You must reboot the switch for the change to take effect.

You can use this command to disable web access to the switch. If you are using ExtremeWare Vista for web access, you must create and configure an access profile before you can use this option. You create an access profile using the `create access-profile` command. You configure an access profile using the `config access-profile` command.

Example

The following command disables web access to the switch:

```
disable web
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

enable cpu-dos-protect

```
enable cpu-dos-protect
```

Description

Enables denial of service protection.

Syntax Description

There are no arguments or variables for this command.

Default

Default is disabled.

Usage Guidelines

None.

Example

The following command enables denial of service protection.

```
enable cpu-dos-protect
```

History

This command was first available in ExtremeWare 6.2.2

Platform Availability

This command is available on all platforms.

enable cpu-dos-protect simulated

```
enable cpu-dos-protect simulated
```

Description

Enables simulated denial of service protection.

Syntax Description

There are no arguments or variables for this command.

Default

Default is disabled.

Usage Guidelines

None.

Example

The following command enables simulated denial of service protection.

```
enable cpu-dos-protect simulated
```

History

This command was first available in ExtremeWare 6.2.2

Platform Availability

This command is available on all platforms.

enable dhcp ports vlan

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

Description

Enables DHCP on a specified port in a VLAN.

Syntax Description

portlist	Specifies the ports for which DHCP should be enabled.
vlan_name	Specifies the VLAN on whose ports DHCP should be enabled.

Default

N/A.

Usage Guidelines

None.

Example

The following command enables DHCP for port 9 in VLAN *corp*:

```
enable dhcp ports 9 vlan corp
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

enable netlogin ports

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

Description

Enables network login on a specified port in a VLAN.

Syntax Description

portlist	Specifies the ports for which netlogin should be enabled.
vlan_name	Specifies the VLAN on whose ports netlogin should be enabled.

Default

N/A.

Usage Guidelines

The VLAN you specify must exist and include the specified ports prior to enabling network login.

For campus mode login, the following conditions must be met:

- A DHCP server must be available, and a DHCP range must be configured for the port or ports in the VLAN on which you want to enable network login.
- The switch must be configured as a RADIUS client, and the RADIUS server must be configured to enable the Extreme network login capability.

For ISP mode login, no special conditions are required. A RADIUS server may be used for authentication, but is not required.

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.

Windows authentication is not supported via network login.

Example

The following command configures network login on port 9 in VLAN *corp*:

```
enable netlogin ports 9 vlan corp
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

enable radius

```
enable radius
```

Description

Enables the RADIUS client on the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

When enabled, all web and 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 authorization before it is executed.

Example

The following command enables RADIUS authentication for the switch:

```
enable radius
```

History

This command was first available in ExtremeWare 4.1.

Platform Availability

This command is available on all platforms.

enable radius-accounting

```
enable radius-accounting
```

Description

Enables RADIUS accounting.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

The RADIUS client must also be enabled.

Example

The following command enables RADIUS accounting for the switch:

```
enable radius-accounting
```

History

This command was first available in ExtremeWare 4.1.

Platform Availability

This command is available on all platforms.

enable snmp access

```
enable snmp access
```

Description

Turns on SNMP support for the switch.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

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

Any network manager running SNMP can manage the switch, provided the MIB is installed correctly on the management station. Each network manager provides its own user interface to the management facilities.

Example

The following command enables SNMP support for the switch:

```
enable snmp access
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

enable snmp dot1dTpFdbTable

```
enable snmp dot1dtpfdbtable
```

Description

Enables SNMP GetNext responses for the dot1dTpFdbTable in the BRIDGE-MIB.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

SNMP Get responses are not affected by this command.

To view the configuration of the dot1dTpFdp table on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable / disable state the dot1dTpFdp table.

To restore defaults to all SNMP-related entries, including the SNMP parameters modified using the `enable snmp dot1dtpfdbtable` command, use the `unconfig management` command.

Example

The following command enables the dot1dTPFdb table:

```
enable snmp dot1dtpfdbtable
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all platforms.

enable snmp traps

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

Description

Turns on SNMP trap support.

Syntax Description

portlist	Specifies the ports for which traps should be enabled.
all	Specifies all ports.

Default

Enabled.

Usage Guidelines

An authorized trap receiver can be one or more network management stations on your network. The switch sends SNMP traps to all trap receivers.

Example

The following command enables SNMP trap support on the switch:

```
enable snmp trap
```

History

This command was first available in ExtremeWare 2.0.

This command was modified to include ports in ExtremeWare 6.2.2.

Platform Availability

This command is available on all platforms.

enable sntp-client

```
enable sntp-client
```

Description

Enables the SNTP client.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

SNTP can be used by the switch to update and synchronize its internal clock from a Network Time Protocol (NTP) server. After the SNTP client has been 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 (DST).

Example

The following command enables the SNTP client:

```
enable sntp-client
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

enable ssh2

```
enable ssh2 {access-profile [<access_profile> | none]} {port
<tcp_port_number>}
```

Description

Enables incoming SSH2 Telnet sessions.

Syntax Description

access profile	Specifies an access profile.
none	Cancels a previously configured access profile.
port	Specifies a TCP port number. The default is port 22.

Default

The SSH2 feature is disabled until you obtain a valid security license. If a valid security license is in effect, SSH2 is enabled with no access profile and uses TCP port number 22.

Usage Guidelines

SSH2 enables the encryption of Telnet session data. You must be logged in as an administrator to enable SSH2 Telnet, and you must obtain and enter a Security License Key to enable the SSH2 feature. To obtain a Security License Key, access the Extreme Networks website.

You can specify a list of predefined clients that are allowed SSH2 access to the switch. To do this, you must create an access profile that contains a list of allowed IP addresses. To create an access profile, use the `create access-profile` command. To configure an access profile, use the `config access-profile` command.

Use the `none` option to cancel a previously configured access profile.

Use the `port` option to specify a TCP port number other than the default.

To view the status of SSH2 Telnet sessions on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for SSH2 Telnet sessions.

Example

The following command enables the SSH2 feature, with access allowed based on the access profile *management*:

```
enable ssh2 management
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on all "*i*" series platforms.

enable system-watchdog

```
enable system-watchdog
```

Description

Enables the system watchdog timer.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

The watchdog timer reboots the switch if the CPU becomes trapped in a processing loop. If the watchdog timer is executed, the switch captures information on the cause of the reboot and posts it to the system log.

You must reboot to have this command take effect.

Example

The following command enables the watchdog timer:

```
enable system-watchdog
```

History

This command was first available in ExtremeWare 6.1.9.

Platform Availability

This command is available on all "*i*" series platforms.

enable tacacs

```
enable tacacs
```

Description

Enables TACACS+.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

After they have been enabled, all web and CLI logins are sent to one of the two TACACS+ servers for login name authentication and accounting.

Example

The following command enables TACACS+ user authentication:

```
enable tacacs
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the "*i*" series platforms.

enable tacacs-accounting

```
enable tacacs-accounting
```

Description

Enables TACACS+ accounting.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

If accounting is used, the TACACS+ client must also be enabled.

Example

The following command enables TACACS+ accounting for the switch:

```
enable tacacs-accounting
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the "*i*" series platforms.

enable tacacs-authorization

```
enable tacacs-authorization
```

Description

Enables CLI command authorization.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

When enabled, each command is transmitted to the remote TACACS+ server for authorization before the command is executed.

Example

The following command enables TACACS+ command authorization for the switch:

```
enable tacacs-authorization
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the "*i*" series platforms.

enable telnet

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

Description

Enables Telnet access to the switch.

Syntax Description

access profile	Specifies an access profile. (6.0, 6.1)
none	Cancels a previously configured access profile. (6.0, 6.1)
port	Specifies a TCP port number. (6.0, 6.1)

Default

Telnet is enabled with no access profile and uses TCP port number 23.

Usage Guidelines

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

If you are using IP without a BOOTP server, you must enter IP parameters for the switch for the Telnet software to communicate with the device. To assign IP parameters to the switch, you must:

- 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.

For version 6.0 and higher:

- Use an access profile to restrict Telnet access. An access profile permits or denies a named list of IP addresses and subnet masks. You must create and configure an access profile before you can use this option. You create an access profile using the `create access-profile` command. You configure an access profile using the `config access-profile` command.
- Use the `none` option to cancel a previously configured access-profile.
- Use the `port` option to specify a TCP port number.

Example

The following command applies the access profile managers to Telnet:

```
enable telnet access-profile managers
```

History

This command was first available in ExtremeWare 2.0.

Support for the `access profile`, `none`, and `port` parameters was introduced in ExtremeWare 6.0.

Platform Availability

This command is available on all platforms.

enable web

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

Description

Enables ExtremeWare Vista web access to the switch.

Syntax Description

access profile	Specifies an access profile. (6.0, 6.1)
none	Cancels a previously configured access profile. (6.0, 6.1)
port	Specifies a TCP port number. (6.0, 6.1)

Default

Enabled, using TCP port 80.

Usage Guidelines

You must reboot the switch for changes to take effect.

For version 6.0 and higher:

- By default, web access is enabled with no access profile and uses TCP port number 80.
- Use an access profile to restrict ExtremeWare Vista web access. An access profile permits or denies a named list of IP addresses and subnet masks. You must create and configure an access profile before you can use this option. You create an access profile using the `create access-profile` command. You configure an access profile using the `config access-profile` command. Apply an access profile only when ExtremeWare Vista is enabled.
- Use the none option to cancel a previously configured access-profile.
- Use the `port` option to specify a TCP port number.

Example

The following command applies the access profile administrators to the web:

```
enable web access-profile administrators
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 6.0 to include the access profile and port options.

Platform Availability

This command is available on all platforms.

exit

```
exit
```

Description

Logs out the session of a current user for CLI or Telnet.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to log out of a CLI or Telnet session.

Example

The following command logs out the session of a current user for CLI or Telnet:

```
exit
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on Summit switches.

logout

```
logout
```

Description

Logs out the session of a current user for CLI or Telnet.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to log out of a CLI or Telnet session.

Example

The following command logs out the session of a current user for CLI or Telnet:

```
logout
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

quit

```
quit
```

Description

Logs out the session of a current user for CLI or Telnet.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to log out of a CLI or Telnet session.

Example

The following command logs out the session of a current user for CLI or Telnet:

```
quit
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

scp2

```
scp2 {cipher [3des | blowfish]} {port <portnum>} {debug <debug_level>}
<user>@ [<hostname> | <ipaddress>] :<remote_file> [configuration
{incremental} | image [primary | secondary] | bootrom]
```

Description

Copies a file from a remote system to the switch using SCP2.

Syntax Description

3des	Specifies that the 3des cipher should be used for encryption. This is the default.
blowfish	Specifies that the blowfish cipher should be used for encryption.
portnum	Specifies the TCP port number to be used for communicating with the SSH2 client. Default is port 22.
debug_level	Specifes a debug level. Default is 0
user	Specifies a login name for the remote host. .
host	Specifies the name of the remote host
ipaddress	Specifies the IP address of the remote host
remote file	Specifes the name of the remote file to be copied to the switch.
configuration	Specifies that the copied file is a switch configuration file. It the incremental option is not specified, it replaces the current switch configuration.
incremental	Specifies that the copied file should be handled like an incremental configuration download (only the commands in the file are executed).
image	Specifies that the copied file is an ExtremeWare image.
primary	Specifies that the image should be placed in the primary image area
secondary	Specifies that the image should be placed in the secondary image area
bootrom	Specifies that the copied file is a bootrom image.

Default

N / A.

Usage Guidelines

You must be running a securtiy-enabled version of ExtremeWare 6.2.1 (which is under Export Control) in order to use the SCP2 command.

SSH2 does not need to be enabled on the switch in order to use this command. (SSH2 is enabled by default if you are running a security-enabled vesion of ExtremeWare).

This command logs into the remote host as <user> and accesses the file <remote_file>. You will be prompted for a password from the remote host, if required.

Example

The following command copies a configuration file from the file *configpart1.save* on host *system1* to the switch as an incremental configuration:

```
scp2 admin@system1:configpart1.save configuration incremental
```

History

This command was first available in ExtremeWare 6.2.1

Platform Availability

This command is available on all "*i*" series platforms.

scp2 configuration

```
scp2 {cipher [3des | blowfish]} {port <portnum>} {debug <debug_level>}
configuration <user>@ [<hostname> | <ipaddress>]:<remote_file>
```

Description

Copies the configuration file from the switch to a remote system using SCP2.

Syntax Description

3des	Specifies that the 3des cipher should be used for encryption. This is the default.
blowfish	Specifies that the blowfish cipher should be used for encryption.
portnum	Specifies the TCP port number to be used for communicating with the SSH2 client. Default is port 22.
debug_level	Specifes a debug level. Default is 0
user	Specifies a login name for the remote host. .
host	Specifies the name of the remote host
ipaddress	Specifies the IP address of the remote host
remote file	Specifes the name of the file to be created on the remote host.

Default

N / A.

Usage Guidelines

You must be running a securtiy-enabled version of ExtremeWare 6.2.1 (which is under Export Control) in order to use the SCP2 command.

SSH2 does not need to be enabled on the switch in order to use this command. (SSH2 is enabled by default if you are running a security-enabled version of ExtremeWare).

This command logs into the remote host as `<user>` and creates the file `<remote_file>`.

Example

The following command copies the switch configuration and saves it as file c*onfig1.save* on host *system1*:

```
scp2 configuration admin@system1:config1.save
```

History

This command was first available in ExtremeWare 6.2.1.

Platform Availability

This command is available on all *"i"* series platforms.

show cpu-dos-protect

```
show cpu-dos-protect
```

Description

Displays the status of denial of service protection.

Syntax Description

There are no arguments or variables for this command.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the status of denial of service protection.

```
show cpu-dos-protect
```

History

This command was first available in ExtremeWare 6.2.2

Platform Availability

This command is available on all platforms.

show management

```
show management
```

Description

Displays the SNMP settings configured on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines:

The following show management output is displayed:

- Enable/disable state for Telnet, SNMP, and web access
- SNMP community strings
- Authorized SNMP station list
- SNMP trap receiver list
- Login statistics

For version 4.0 and higher, the following show management output is also displayed:

- Enable/disable state for SSH2 and access profile information
- RMON polling configuration

Example

The following command displays configured SNMP settings on the switch:

```
show management
```

Following is the output from this command:

```
CLI idle timeouts:              disabled
CLI Paging:                     enabled
CLI configuration logging:      enabled
Telnet access:                  enabled tcp port: 23
Web access:                     enabled tcp port: 80
SSH Access:                     key invalid, disabled tcp port: 22
UDP Echo Server:                disabled udp port: 7
SNMP Access:                    enabled
SNMP Read Only Communities:     rykfcb
Total Read Only Communities:    1
SNMP Read Write Communities:    r~`|kug
Total Read Write Communities:   1
SNMP dot1dTpFdbTable:           disabled
RMON polling:                   disabled
```

```
SNMP Traps:                     enabled
SNMP TrapReceivers:
    Destination           Community             Source IP Address   Flags
   10.0.5.117 /10550      ST.167773557.10550                        E----------
   111.111.111.111/162    ThisIsATestComm       10.203.0.149        S----------
   222.111.111.111/162    public                                    E----------

Flags : S=Standard Mode, E=Enhanced Mode

Link Up/Link Down traps enabled on ports:
10   11   12   13   14   15   16   17
18   19   20   21   22   23   24   25
26   27   28   29   30   31   32   33

SNMP stats:     inPkts 10       outPkts 20      errors 0        authErrors 0
                Gets 8             GetNexts 2     Sets 0
SNMP traps:     sent 10 authTraps enabled
```

History

This command was first available in ExtremeWare 2.0.

Support for the SSH2 state, access profile information and RMON polling configuration was introduced in ExtremeWare 4.0.

Additional information on traps configured per port was added in ExtremeWare 6.2.2.

Platform Availability

This command is available on all platforms.

show netlogin info

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

Description

Disables network login on a specified port in a VLAN.

Syntax Description

portlist	Specifies the ports for which netlogin should be disabled.
vlan_name	Specifies the VLAN on whose ports netlogin should be disabled.

Default

N/A.

Usage Guidelines

The information reported by this command is the following:

- The port and VLAN for which the information is displayed.
- The port state: Authenticated or Not Authenticated.
- The temporary IP assigned, if known.
- The DHCP state: Enabled or Not Enabled.
- The user name, if known.
- The MAC address of the attached client, if known.

Example

The following command shows network login information for port 9 on VLAN *corp*:

```
show netlogin info ports 9 vlan corp
```

The results of this command are as follows:

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

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

show odometer

```
show odometer
```

Description

Displays a counter for each component of a switch that shows how long it has been functioning since it was manufactured.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The output from this command displays how long each individual component in the whole switch has been functioning since it is manufactured. This odometer counter will be kept in the EEPROM of each monitored component. This means that even when the component is plugged into different chassis, the odometer counter will be available in the new switch chassis. The following components are monitored by the odometer:

- For the Black Diamond—MSM and I/O modules
- For the Alpine—SMM, I/O slots, and power supplies
- For stackable switches—the CPU

Example

The following command displays how long each component of a switch has been functioning since it's manufacture date:

```
show odometer
```

The results of this command are as follows:

```
MSM64:7 # show odometers
                             Service
Field Replaceable Units      Days       First Recorded Start Date
---------------------------- ------------ -------------------------
Backplane:
MSM-A:           MSM64i           1    Mon Mar 4 10:28:28 2002
MSM-B:           MSM64i           1    Mon Mar 4 10:28:28 2002
Slot 1:          F48Ti            1    Mon Mar 4 10:28:28 2002
Slot 2:          F48Ti            1    Mon Mar 4 10:28:28 2002
Slot 3:          F48Ti            1    Mon Mar 4 10:28:28 2002
Slot 4:          F48Ti            1    Mon Mar 4 10:28:28 2002
Slot 5:          F48Ti            1    Mon Mar 4 10:28:28 2002
Slot 6:          F48Ti            1    Mon Mar 4 10:28:28 2002
Slot 7:          F48Ti            1    Mon Mar 4 10:28:28 2002
Slot 8:          F48Ti            1    Mon Mar 4 10:28:28 2002
```

History

This command was first available in ExtremeWare 6.2.1.

Platform Availability

This command is available on all platforms.

show radius

```
show radius
```

Description

Displays the current RADIUS client configuration and statistics.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The output from this command displays the status of the RADIUS and RADIUS accounting (enabled or disabled) and the primary and secondary servers for RADIUS and RADIUS accounting:

Example

The following command displays the current RADIUS client configuration and statistics:

```
show radius
```

History

This command was first available in ExtremeWare 4.1.

Platform Availability

This command is available on all platforms.

show radius-accounting

```
show radius-accounting
```

Description

Displays the current RADIUS accounting client configuration and statistics.

Syntax Description

This command has no arguments or variables.

Default

N / A.

Usage Guidelines

The output from this command displays information about the status and configuration of RADIUS accounting

Example

The following command displays RADIUS accounting client configuration and statistics:

```
show radius-accounting
```

History

This command was first available in ExtremeWare 4.1.

Platform Availability

This command is available on all platforms.

show session

```
show session
```

Description

Displays the currently active Telnet, console, and web sessions communicating with the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The `show session` command displays the username and IP address of the incoming Telnet session, whether a console session is currently active, and the login time.

The following table displays the `show session` command field definitions.

Table 7: Show Command Field Definitions

Field	Definition
#	Indicates session number.
Login Time	Indicates login time of session.
User	Indicates the user logged in for each session.
Type	Indicates the type of session.
Auth	Indicates how the user is logged in.
CLI Auth	Indicates the type of authentication (RADIUS and TACAS) if enabled.
Location	Indicates the location (IP address) from which the user logged in.

Example

The following command displays the active sessions on the switch:

```
show session
```

It produces output similar to the following:

```
# Login Time                 User     Type     Auth     CLI Auth Location
================================================================================
     0 Tue Feb 19 18:08:42 2002 admin    console  local    disabled serial
     5 Thu Feb 21 19:09:48 2002 admin    http     local    disabled 10.0.4.76
* 1028 Thu Feb 21 18:56:40 2002 admin    telnet   local    disabled 10.0.4.19
```

History

This command was first available in ExtremeWare 2.0.

Support for the CLI Auth command field definition was introduced in ExtremeWare 6.0.

Support for the Auth command field definition was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

show sntp-client

```
show sntp-client
```

Description

Displays the DNS configuration.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Displays configuration and statistics information of SNTP client.

Example

The following command displays the DNS configuration:

```
show sntp-client
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

show tacacs

```
show tacacs
```

Description

Displays the current TACACS+ configuration and statistics.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays TACACS+ client configuration and statistics:

```
show tacacs
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the "*i*" series platforms.

show tacacs-accounting

```
show tacacs-accounting
```

Description

Displays the current TACACS+ accounting client configuration and statistics.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None:

Example

The following command displays TACACS+ accounting client configuration and statistics:

```
show tacacs-accounting
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the "*i*" series platforms.

ssh2

```
ssh2 {cipher [3des | blowfish]} {port <portnum>} {compression [on | off]}
{user <username>} {debug <debug_level>} {<username>@} [<host> |
<ipaddress>] {<remote command>}
```

Description

Transmits a command to a remote system using an SSH2 connection.

Syntax Description

3des	Specifies that the 3des cipher should be used for encryption. This is the default.
blowfish	Specifies that the blowfish cipher should be used for encryption.
portnum	Specifies the TCP port number to be used for communicating with the SSH2 client. Default is port 22.
compression	on specifies that data is to be compressed. `off` specifies that compression is not to be used. Default is `off`.
username	Specifies a login name for the remote host, as an alternate to the user@host parameter.
debug_level	Specifes a debug level. Default is 0
username	Specifies a login name for the remote host. May be omitted if it is the same as the username on the switch.
host	Specifies the name of the remote host
ipaddress	Specifies the IP address of the remote host
remote command	Specifes a command to be passed to the remote system for execution. Remote commands are not supported on switches. This option is only valid if the remote system is a system, such as a UNIX workstation, that can accept remote commands.

Default

N / A.

Usage Guidelines

You must be running a securtiy-enabled version of ExtremeWare 6.2.1 (which is under Export Control) in order to use the SSH2 client command.

SSH2 does not need to be enabled on the switch in order to use this command. (SSH2 is enabled by default if you are running a scurity-enabled version of ExtremeWare).

Typically this command is used to establish a secure session to a remote switch. You will be prompted for your password. Once you have logged in successfully, all ExtremeWare commands you enter will be exectued on the remote switch. When you terminate the remote session, commands will then resume being executed on the original switch.

The remote command option cannot be used with Extreme Networks switches. If you include a remote command, you will receive an error message.

Example

The following command establishes an SSH2 session on switch engineering1:

```
ssh2 admin@engineering1
```

The following command establishes an SSH2 session with the switch summit48i over TCP port 2050 with compression enabled:

```
ssh2 port 2050 compression on admin@summit48i
```

History

This command was first available in ExtremeWare 6.2.1

Platform Availability

This command is available on all *"i"* series platforms.

telnet

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

Description

Allows you to Telnet from the current command-line interface session to another host.

Syntax Description

ipaddress	Specifies the IP address of the host.
hostname	Specifies the name of the host. (4.x and higher)
port_number	Specifies a TCP port number. (4.x and higher)

Default

Enabled. If the TCP port number is not specified, the Telnet session defaults to port 23.

Usage Guidelines

Only VT100 emulation is supported.

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

You need to configure the switch IP parameters.

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

Before you can start a Telnet session, you need to configure the switch IP parameters. To open a Telnet connection, you must specify the host IP address or the host name of the device you wish to manage. Check the user manual supplied with the Telnet facility if you are unsure of how to do this.

To view the status of Telnet on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for Telnet.

For version 4.x and higher:

- You must configure DNS in order to use the `hostname` option.

For version 2.0:

- The `hostname` parameter is not available.

Example

The following command configures Telnet communication with a host at IP address 123.45.67.8:

```
telnet 123.45.67.8
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 4.x to support the `hostname` and `port number` parameters.

Platform Availability

This command is available on all platforms.

unconfig management

```
unconfig management
```

Description

Restores default values to all SNMP-related entries.

Syntax Description

This command has no arguments or variables.

Default

N/A

Usage Guidelines

None.

Example

The following command restores default values to all SNMP-related entries on the switch:

```
unconfig management
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

unconfig radius

```
unconfig radius {server [primary | secondary]}
```

Description

Unconfigures the RADIUS client configuration.

Syntax Description

primary	Unconfigures the primary RADIUS server.
secondary	Unconfigures the secondary RADIUS server.

Default

Unconfigures both primary and secondary servers.

Usage Guidelines

None.

Example

The following command unconfigures the secondary RADIUS server for the client:

```
unconfig radius server secondary
```

History

This command was first available in ExtremeWare 4.1.

Platform Availability

This command is available on all platforms.

unconfig radius-accounting

```
unconfig radius-accounting {server [primary | secondary]}
```

Description

Unconfigures the RADIUS accounting client configuration.

Syntax Description

primary	Unconfigures the primary RADIUS accounting server.
secondary	Unconfigures the secondary RADIUS accounting server.

Default

Unconfigures both the primary and secondary accounting servers.

Usage Guidelines

None.

Example

The following command unconfigures the secondary RADIUS accounting server for the client:

```
unconfig radius-accounting server secondary
```

History

This command was first available in ExtremeWare 4.1.

Platform Availability

This command is available on all platforms.

unconfig tacacs

```
unconfig tacacs {server [primary | secondary]}
```

Description

Unconfigures the TACACS+ client configuration.

Syntax Description

primary	Unconfigures the primary TACACS+ server.
secondary	Unconfigures the secondary TACACS+ server.

Default

Unconfigures both the primary and secondary TACACS+ servers.

Usage Guidelines

None.

Example

The following command unconfigures all TACACS+ servers for the client:

```
unconfig tacacs
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the "*i*" series platforms.

unconfig tacacs-accounting

```
unconfig tacacs-accounting {server [primary | secondary]}
```

Description

Unconfigures the TACACS+ accounting client configuration.

Syntax Description

primary	Unconfigures the primary TACACS+ accounting server.
secondary	Unconfigures the secondary TACACS+ accounting server.

Default

Unconfigures both the primary and secondary TACACS+ accounting servers.

Usage Guidelines

None.

Example

The following command unconfigures all TACACS+ accounting servers for the client:

```
unconfig tacacs-accounting
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the "*i*" series platforms.

4 Commands for Configuring Slots and Ports on a Switch

This chapter describes:

- Commands related to enabling, disabling, and configuring individual ports
- Commands related to configuring port speed (Fast Ethernet ports only) and half- or full-duplex mode
- Commands related to creating load-sharing groups on multiple ports
- Commands related to displaying port statistics

By default, all ports on the switch are enabled. After you configure the ports to your specific needs, you can select which ports are enabled or disabled.

Fast Ethernet ports can connect to either 10Base-T or 100Base-T networks. By default, the ports autonegotiate (automatically determine) the 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.

The switch comes configured to use autonegotiation to determine the port speed and duplex setting for each port. You can select to manually configure the duplex setting and the speed of 10/100 Mbps ports, and you can manually configure the duplex setting on Gigabit Ethernet ports.

All ports on the switch can be configured for half-duplex or full-duplex operation. The ports are configured to autonegotiate the duplex setting, but you can manually configure the duplex setting for your specific needs.

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.

Load sharing with Extreme Network switches allows you to increase bandwidth and resilience between switches 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. The algorithm also guarantees 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.

Load sharing is most useful in cases where the traffic transmitted from the switch to the load-sharing group is sourced from an equal or greater number of ports on the switch. For example, traffic

transmitted to a two-port load-sharing group should originate from a minimum of two other ports on the same switch.

You can view port status on the switch using the `show ports` commands. These commands, when used with specific keywords and parameters, allow you to view various issues such as real-time collision statistics, link speed, flow control, and packet size.

Commands that require you to enter one or more port numbers use the parameter `<portlist>` in the syntax. On a modular switch, a `<portlist>` can be a list of slots and ports. On a stand-alone switch, a `<portlist>` can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

clear slot

```
clear slot <slot>
```

Description

Clears a slot of a previously assigned I/O module type.

For MPLS modules, clears a slot of a previously assigned MPLS module.

Syntax Description

slot	Specifies a modular switch slot number.

Default

N/A.

Usage Guidelines

All configuration information related to the slot and the ports on the module is erased. If a module is present when you issue this command, the module is reset to default settings.

If a slot is configured for one type of module, and a different type of module is inserted, the inserted module is put into a mismatch state (where the inserted module does not match the configured slot), and is not brought online. To use the new module type in a slot, the slot configuration must be cleared or configured for the new module type. Use the `config slot` command to configure the slot.

For version 6.0 and higher:

- This command is available on modular switches.

For version 4.0:

- This command is available on BlackDiamond switches only.

Example

The following command clears slot 2 of a previously assigned I/O module type:

```
clear slot 2
```

The following command clears slot 3 of a previously assigned MPLS module:

```
clear slot 3
```

History

This command was first available in ExtremeWare 4.0.

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12 to support MPLS modules.

Platform Availability

This command is available on modular switches only.

config ip-mtu vlan

```
config ip-mtu <size> vlan <vlan name>
```

Description

Sets the maximum transmission unit (MTU) for the VLAN.

For MPLS modules, configures the IP MTU size.

Syntax Description

size	Specifies the ip-mtu value. The ip-mtu value can be either 1500 or 9216.
vlan name	Specifies a VLAN name.

Default

The default IP MTU size is 1500.

Usage Guidelines

If you enter a value other than 1500 or 9216, the switch will use the value 9216.

Use this command to enable jumbo frame support or for IP fragmentation with jumbo frames. Jumbo frames are Ethernet frames that are larger than 1523 bytes, including 4 bytes used for CRC. Both endstations involved in the transfer must be capable of supporting jumbo frames. The switch does not perform IP fragmentation or participate in MTU negotiation on behalf of devices that do not support jumbo frames.

When enabling jumbo frames and setting the MTU size for the VLAN, keep in mind that some network interface cards (NICs) have a configured maximum MTU size that does not include the additional 4 bytes of CRC. Ensure that the NIC maximum MTU is at or below the maximum MTU size configured on the switch. Frames that are larger than the MTU size configured on the switch are dropped at the ingress port.

If you use IP fragmentation with jumbo frames and you want to set the MTU size greater than 1500, all ports in the VLAN must have jumbo frames enabled.

For MPLS modules:

Fragmentation is based on either the minimum value of the configured MPLS IP MTU size or the configured IP MTU size for the egress VLAN.

Example

The following command sets the MTU size to 1500 for VLAN *sales*:

```
config ip-mtu 1500 vlan sales
```

The following command increases the MTU size on the MPLS VLANs to accommodate the MPLS shim header:

```
config ip-mtu 1550 vlan vlan1
```

History

This command was first available in ExtremeWare 6.2.

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12 to support MPLS modules.

Platform Availability

This command is available on "*i*" series platforms.

config jumbo-frame size

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

Description

Sets the maximum jumbo frame size for the switch chassis.

Syntax Description

number	Specifies a maximum transmission unit (MTU) size for a jumbo frame.

Default

The default setting is 9216.

Usage Guidelines

Jumbo frames are used between endstations that support larger frame sizes for more efficient transfers of bulk data. Both endstations involved in the transfer must be capable of supporting jumbo frames.

The `number` keyword describes the maximum jumbo frame size "on the wire," and includes 4 bytes of cyclic redundancy check (CRC) plus another 4 bytes if 802.1Q tagging is being used.

To enable jumbo frame support, you must configure the maximum transmission unit (MTU) size of a jumbo frame that will be allowed by the switch.

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

For MPLS modules:

You should enable jumbo frame support on the ports that are members of an MPLS VLAN. The jumbo frame size should be set to accommodate the addition of a maximally-sized label stack. For example, a jumbo frame size of at least 1530 bytes is needed to support a two-level label stack on a tagged Ethernet port and a jumbo frame size of at least 1548 bytes is needed to support a TLS encapsulated MPLS frame.

The MPLS module supports the MTU size configured using the `config jumbo-frame size` command.

For version 6.1 and higher:

- The jumbo_frame_mtu range is between 1523 through 9216.

For version 6.0:

- The jumbo_frame_mtu range is between 1522 through 9216.

Example

The following command configures the maximum MTU size of a jumbo frame size to 5500:

```
config jumbo-frame size 5500
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

config mirroring add

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

Description

Adds a particular mirroring filter definition on the switch.

Syntax Description

mac_address	Specifies a MAC address. (Supported in versions 2.0 - 4x only)
vlan name	Specifies a VLAN name.
portnumber	Specifies a port or slot and port.

Default

N/A.

Usage Guidelines

On a modular switch, <portnumber> will be a slot and port in the form <slot>:<port>. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

You must enable port-mirroring using the enable mirroring command before you can configure the mirroring filter definitions.

Up to eight mirroring definitions can be added. You can mirror traffic from a VLAN, a physical port, or a specific VLAN/port combination.

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

- Physical port — All data that traverses the port, regardless of VLAN configuration, is copied to the monitor port.
- VLAN — All data to and from a particular VLAN, regardless of the physical port configuration, is copied to the monitor port.
- Virtual port — All data specific to a VLAN on a specific port is copied to the monitor port.

Up to eight mirroring filters and one monitor port can be configured on the switch. Once a port is specified as a monitor port, it cannot be used for any other function. Frames that contain errors are not mirrored.

For version 2.0 and 4.0:

In addition to the physical port, VLAN, and virtual port, the traffic filter can be defined based on the following criteria:

- MAC source address/destination address — All data sent to or received from a particular source or destination MAC address is copied to the monitor port.

For MAC mirroring to work correctly, the MAC address must already be present in the forwarding database (FDB). You need to enable and configure FDB for MAC mirroring to work correctly. See "FDB Commands" for more details.

Example

The following example sends all traffic coming into or out of a stand-alone switch on port 1 and the VLAN *default* to the mirror port:

```
config mirroring add ports 1 vlan default
```

The following example sends all traffic coming into or out of a modular switch on slot 3, port 2 and the VLAN *default* to the mirror port:

```
config mirroring add ports 3:2 vlan default
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 4.0 to support modular switches.

This command was modified in ExtremeWare 6.0 to discontinue support for the MAC address parameter.

Platform Availability

This command is available on all platforms.

config mirroring delete

```
config mirroring delete [<mac_address> | vlan <vlan name> {ports
<portnumber>} | ports <portnumber> {vlan <vlan name>}]
```

Description

Deletes a particular mirroring filter definition on the switch.

Syntax Description

mac_address	Specifies a MAC address. (Supported in versions 4.0 and 6.0 only)
vlan name	Specifies a VLAN name.
portnumber	Specifies a port or slot and port.

Default

N / A.

Usage Guidelines

On a modular switch, <portnumber> must be a slot and port in the form <slot>:<port>. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

For version 6.0:

- No longer supports using a MAC address to specify mirroring.

Example

The following example deletes the mirroring filter on a stand-alone switch defined for port 1 on VLAN *default*:

```
config mirroring delete ports 1 vlan default
```

The following example deletes the mirroring filter on a modular switch defined for slot 3, port 2 on VLAN default:

```
config mirroring add ports 3:2 vlan default
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 4.0 to support modular switches.

This command was modified in ExtremeWare 6.0 to discontinue support for the MAC address parameters.

Platform Availability

This command is available on all platforms.

config ports

```
config ports [<portlist> vlan <name> | all] [limit-learning <number> |
lock-learning | unlimited-learning | unlock-learning]
```

Description

Configures virtual ports for limited or locked MAC address learning.

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
all	Specifies that all virtual ports should be configured as indicated.
name	Specifies the name of the VLAN
limit-learning <number>	Specifies a limit on the number of MAC addresses that can be dynamically learned on the specified ports.
lock-learning	Specifies that the current FDB entries for the specified ports should be made permanent static, and no additional learning should be allowed.
unlimited-learning	Specifies that there should not be a limit on MAC addresses that can be learned.
unlock-learning	Specifies that the port should be unlocked (allow unlimited, dynamic learning).

Default

Unlimited, unlocked learning.

Usage Guidelines

On a modular switch, `<portlist>` can be a list of slots and ports. On a stand-alone switch, `<portlist>` can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

Limited learning. The limited learning feature allows you to limit the number of dynamically-learned MAC addresses per VLAN. When the learned limit is reached, all new source MAC addresses are blackholed at both the ingress and egress points. This prevent these MAC addresses from learning and responding to Internet control message protocol (ICMP) and address resolution protocol (ARP) packets.

If the limit you configure is greater than the current number of learned entries, all the current learned entries are purged.

Dynamically learned entries still get aged, and can be cleared. If entries are cleared or aged out after the learning limit has been reached, new entries will then be able to be learned until the limit is reached again.

Permanent static and permanent dynamic entries can still be added and deleted using the `create fdbentry` and `delete fdbentry` commands. These override any dynamically learned entries.

For ports that have a learning limit in place, the following traffic will still flow to the port:

- Packets destined for permanent MACs and other non-blackholed MACs

- Broadcast traffic
- EDP traffic

Traffic from the permanent MAC and any other non-blackholed MACs will still flow from the virtual port.

If you configure a MAC address limit on VLANS that have ESRP enabled, you should add an additional back-to-back link (that has no MAC address limit on these ports) between the ESRP-enabled switches. Doing so prevents ESRP PDU from being dropped due to MAC address limit settings.

Port lockdown. The port lockdown feature allows you to prevent any additional learning on the virtual port, keeping existing learned entries intact. This is equivalent to making the dynamically-learned entries permanent static, and setting the learning limit to zero. All new source MAC addresses are blackholed.

Locked entries do not get aged, but can be deleted like any other permanent FDB entries. The maximum number of permanent lockdown entries is 1024. Any FDB entries above will be flushed and blackholed during lockdown.

For ports that have lockdown in effect, the following traffic will still flow to the port:

- Packets destined for the permanent MAC and other non-blackholed MACs
- Broadcast traffic
- EDP traffic

Traffic from the permanent MAC will still flow from the virtual port.

Once the port is locked down, all the entries become permanent and will be saved across reboot. When you remove the lockdown using the unlock-learning option, the learning-limit is reset to unlimited, and all associated entries in the FDB are flushed.

To verify the MAC security configuration for the specified VLAN or ports, use the following commands:

```
show vlan <name> security
show ports <portlist> info detail
```

Example

The following command limits the number of MAC addresses that can be learned on ports 1, 2, 3, and 6 in a VLAN named *accounting*, to 128 addresses:

```
config ports 1, 2, 3, 6 vlan accounting learning-limit 128
```

The following command locks ports 4 and 5 of VLAN *accounting*, converting any FDB entries to static entries, and prevents any additional address learning on these ports:

```
config ports 4,5 vlan accounting lock-learning
```

The following command removes the learning limit from the specified ports:

```
config ports 1, 2, vlan accounting unlimited-learning
```

The following command unlocks the FDB entries for the specified ports:

```
config ports 4,5 vlan accounting unlock-learning
```

History

This command was first available in ExtremeWare 6.2.1

Platform Availability

This command is available on the "*i*" series platforms.

config ports auto off

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

Description

Manually configures port speed and duplex setting configuration on one or more ports on a switch.

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
all	Specifies all configured ports on the switch. (6.1 and higher) See “Usage Guidelines” for more information.
mgmt	Specifies the management port. Supported only for switches that provide a management port.
speed [10]	Specifies 10 Mbps ports.
speed [100]	Specifies 100 Mbps ports.
speed [1000]	Specifies 1000 Mbps ports. (6.1 and higher)
duplex [half]	Specifies half duplex; transmitting and receiving data one direction at a time.
duplex [full]	Specifies full duplex; transmitting and receiving data at the same time.

Default

Auto on.

Usage Guidelines

On a modular switch, <portlist> can be a list of slots and ports. On a stand-alone switch, <portlist> can be one or more port numbers. For a detailed explanation of port specification, see “Modular Switch Numerical Ranges” or “Stand-alone Switch Numerical Ranges” in Chapter 1.

You can manually configure the duplex setting and the speed of 10/100 Mbps ports, and you can manually configure the duplex setting on Gigabit Ethernet ports.

Fast Ethernet 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 ports autonegotiate the duplex setting.

In certain interoperability situations, it is necessary to turn autonegotiation off on a Gigabit Ethernet port. Even though a Gigabit Ethernet port runs only at full duplex and gigabit speeds, the command that turns off autonegotiation must still include the duplex setting.

Gigabit ethernet ports support flow control only when autonegotiation is turned on. When autonegotiation is turned off, flow control is not supported.

For version 6.1:

- The `all` parameter specifies all ports on the switch.
- The `1000` parameter specifies 1000 Mbps ports.

Example

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
```

The following example turns autonegotiation off for slot 2, port 1 on a modular switch:

```
config ports 2:1 auto off duplex full
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 4.0 to support modular switches.

This command was modified in ExtremeWare 6.1 to support the `all` parameter.

Platform Availability

This command is available on all platforms.

config ports auto on

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

Description

Enables autonegotiation for the particular port type.

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
mgmt	Specifies the management port. Supported only for switches that provide a management port.
all	Specifies all configured ports on the switch. (6.1 and higher) See “Usage Guidelines” for more information.

Default

Auto on.

Usage Guidelines

On a modular switch, <portlist> can be a list of slots and ports. On a stand-alone switch, <portlist> can be one or more port numbers. For a detailed explanation of port specification, see “Modular Switch Numerical Ranges” or “Stand-alone Switch Numerical Ranges” in Chapter 1.

The type of ports enabled for autonegotiation are 802.3u for 10/100 Mbps ports or 802.3z for Gigabit Ethernet ports.

Flow control is supported on Gigabit Ethernet ports only. 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.

For version 6.1:

- The `all` parameter specifies all ports on the switch.

Example

The following command configures the switch to autonegotiate for ports 4 and 6 on a stand-alone switch:

```
config ports 4,6 auto on
```

The following command configures the switch to autonegotiate for slot 1, ports 2 and 4 on a modular switch:

```
config ports 1:2, 1:4 auto on
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 6.1 to support the `all` parameter.

This command was modified in ExtremeWare 4.0 to support modular switches.

Platform Availability

This command is available on all platforms.

config ports display-string

```
config ports [<portlist> | mgmt] display-string <alphanumeric string>
```

Description

Configures a user-defined string for a port or group of ports.

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
mgmt	Specifies the management port. Supported only for switches that provide a management port.
alphanumeric string	Specifies a user-defined display string.

Default

N / A.

Usage Guidelines

On a modular switch, <portlist> can be a list of slots and ports. On a stand-alone switch, <portlist> can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

The display string can be up to 16 characters. Display strings do not need to be unique for each port—you can assign the same string to multiple ports. For example, you could give all the ports that connected to a particular department a common display string.

The string is displayed in certain commands such as the `show ports info` command.

Example

The following command configures the user-defined string *corporate* for port 1 on a stand-alone switch:

```
config ports 1 display-string corporate
```

The following command configures the user-defined string *corporate* for ports 3, 4, and 5 on slot 1 on a modular switch:

```
config ports 1:3-5 display-string corporate
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

config ports redundant

```
config ports [<portlist> | <portid> | mgmt] redundant [<portlist> |
<portid>]
```

Description

Configures a software-controlled redundant port.

Syntax Description

portlist	Specifies one or more primary ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
portid	Specifies a primary port using the display string configured for the port. If this option is used to identify the port, the redundant port must also be specified using a port id (display string).
mgmt	Specifies the management port as the primary port. Supported only for switches that provide a management port.
portlist	Specifies one or more redundant ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
portid	Specifies a redundant port using the display string configured for the port. This option may be used to identify the redundant port of the primary port was specified using the port id (display string).

Default

N / A.

Usage Guidelines

On a modular switch, `<portlist>` can be a list of slots and ports. On a stand-alone switch, `<portlist>` can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

The first port list specifies the primary ports. The second port list specifies the redundant ports.

A software-controlled redundant port is configured to backup a specified primary port. The redundant port tracks the link state of the associated primary port, and if the link on the primary port fails, the redundant port establishes a link and becomes active. You can back up a specified Ethernet port with a redundant, dedicated Ethernet port. You can also back up a load-shared group of Ethernet ports with a set of load-shared redundant Ethernet ports. If a link in the active load-shared group fails, the entire group fails over to the redundant group.

The following criteria must be considered when configuring a software-controlled redundant port:

- You must manually configure the primary and redundant ports identically in terms of VLANs, QoS settings, access lists, and so on.
- Auto-negotiation must be enabled on both the primary and redundant port.
- You cannot configure hardware redundant ports (such as ports 49 and 50 on a Summit48i) as software controlled redundant ports.
- Software redundant ports are supported on products that use the "*i*" chipset.

- Only one side of the link should be configured as redundant. For example, if ports 1 and 2 are connected between switches A and B, only switch A should be configured with redundant ports.
- Software redundant ports are not supported on 1000BASE-T ports.

Software redundant port only cover failures where both the TX and RX paths fail. If a single strand of fiber is pulled, the software redundant port cannot correctly recover from the failure.

Example

The following command configures a software-controlled redundant port on a stand-alone switch:

```
config ports 3 redundant 4
```

The following command configures a software-controlled redundant port on a modular switch:

```
config ports 1:3 redundant 2:3
```

The following command configures a software-controlled redundant port using the port display strings corp1 and corp5 to identify the ports:

```
config ports corp1 redundant corp5
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config sharing address-based

```
config sharing address-based [L2 | L2_L3 | L2_L3_L4]
```

Description

Configures the part of the packet examined by the switch when selecting the egress port for transmitting load-sharing data.

Syntax Description

L2	Indicates that the switch should examine the MAC source and destination address.
L2-L3	Indicates that the switch should examine the IP source and destination address.
L2-L3-L4	Indicates that the switch should examine the UDP or TCP well-know port number.

Default

N / A.

Usage Guidelines

This feature is available using the address-based load-sharing algorithm only. 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 address, and the TCP port number.
- IPX packets—Uses the source and destination MAC address and IPX identifiers.
- All other packets—Uses the source and destination MAC address.

To verify your configuration, use the `show sharing address-based` command. The show sharing address-based output displays the addressed-based configurations on the switch.

Example

The following example configures the switch to examine the MAC source and destination address:

```
config sharing address-based l2
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the *"i"* series platform and the Alpine 3800 series switch modules.

config slot

```
config slot <slot> module <module name>
```

Description

Configures a slot for a particular I/O module card in a modular switch.

Syntax Description

slot	Specifies the slot number.
module name	Specifies the type of module for which the slot should be configured. The list of modules you can enter will vary depending on the type of switch and version of ExtremeWare you are running. Certain modules are supported only with specific ExtremeWare Technology Releases. The following are some of the modules you may specify for a BlackDiamond switch: f32f — Specifies Fast Ethernet, 32-port, fiber module. f32fi — Specifies Fast Ethernet, 32-port, fiber module, “i” chipset. f32t — Specifies Fast Ethernet, 32-port, copper module. f48t — Specifies Fast Ethernet, 48-port, copper module. f96t —Specifies Fast Ethernet, 96-port, copper module. g4x — Specifies a Gigabit Ethernet, 4-port, copper module. g6x — Specifies a Gigabit Ethernet, 6-port, copper module. g8t — Specifies a Gigabit Ethernet, 8-port, copper module. g8x — Specifies a Gigabit Ethernet, 8-port, copper module. g12sx — Specifies a Gigabit Ethernet, 12-port, fiber module. g12tx — Specifies a Gigabit Ethernet, 12-port, copper module. WDMi — Specifies a Gigabit Ethernet, WAN module. arm — Specified an Accounting and Routing Module (ARM). a3c — Specifies an Asynchronous Transfer Mode (ATM) module. mpls — Specifies a MultiProtocol Label Switching (MPLS) module. p3c — Specifies an OC-3 PoS module. p12c — Specifies an OC-12 PoS module. The following are some of the modules you may specify for an Alpine switch: fm8v — Specifies a VDSL module. fm24t — Specifies a Fast Ethernet, 24-port, copper module. fm24mf — Specifies a Fast Ethernet, 24-port, multi-mode, fiber module. fm24sf — Specifies a Fast Ethernet, 24-port, single mode, fiber module. fm32t — Specifies a Fast Ethernet, 32-port, copper module. gm4s — Specifies Gigabit Ethernet, 4-port, fiber module. gm4t — Specifies a Gigabit Ethernet, 4-port, copper module. gm4x — Specifies Gigabit Ethernet, 4-port, GBIC module. wdmi — Specifies a Gigabit Ethernet WAN module. (6.1 or higher)

	wm4t1 — Specifies a T1 WAN module. (6.1 or higher) wm4e1 — Specifies an E1 WAN module. wm1t3 — Specifies a T3 WAN module.

Default

If a slot has not been configured for a particular type of I/O module, then any type of module is accepted in that slot, and a default port and VLAN configuration is automatically generated.

Usage Guidelines

The `config slot` command displays different module parameters depending on the type of modular switch you are configuring and the version of ExtremeWare running on the switch.

You can also preconfigure the slot before inserting the module card. This allows you to begin configuring the module and ports before installing the card in the chassis.

If a slot has not been configured for a particular type of I/O module, then any type of module is accepted in that slot, and a default port and VLAN configuration is automatically generated. If a slot is configured for one type of module, and a different type of module is inserted, the inserted module is put into a mismatch state, and is not brought online. To use the new module type in a slot, the slot configuration must be cleared or configured for the new module type.

Upon powering up the chassis, or when an I/O module is hot-swapped, ExtremeWare automatically determines the system power budget and protects the BlackDiamond switch from any potential overpower configurations. If power is available, ExtremeWare powers on and initializes the module. When ExtremeWare detects that a module will cause an overpower condition, the module remains powered down, and is not initialized. An entry is made to the system log indicating the condition.

For version 4.0:

- This command is available on BlackDiamond switches only.

Example

The following command configures the slot for a Fast Ethernet, 32-port, copper module:

```
config slot 2 module F32T
```

History

This command was first available in ExtremeWare 4.0.

This command was modified in an ExtremeWare IP Technology Services Release based on v6.1.8b12 to support MPLS modules.

This command was modified in ExtremeWare 6.1 to support the PoS modules and additional Alpine I/O modules.

This command was modified in ExtremeWare 6.0 to support the Alpine and additional BlackDiamond F48T, G8X, and G12X I/O modules.

Platform Availability

This command is available on modular switches only.

disable edp ports

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

Description

Disables the Extreme Discovery Protocol (EDP) on one or more ports.

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
all	Specifies all ports on the switch. (6.1) See “Usage Guidelines” for more information.

Default

Enabled.

Usage Guidelines

On a modular switch, `<portlist>` can be a list of slots and ports. On a stand-alone switch, `<portlist>` can be one or more port numbers. For a detailed explanation of port specification, see “Modular Switch Numerical Ranges” or “Stand-alone Switch Numerical Ranges” in Chapter 1.

You can use the `disable edp ports` command to disable EDP on one or more ports when you no longer need to locate neighbor Extreme Networks switches.

For version 6.1:

- The `all` parameter specifies all ports on the switch.

For Version 6.0 and higher:

- SummitLink is not supported.

For version 2.0 and 4.0:

- EDP cannot be disabled on a port that has SummitLink enabled, nor on ports that are connected to a Summit Virtual Chassis.

Example

The following command disables EDP on port 4 and port 6 on a stand-alone switch:

```
disable edp ports 4,6
```

The following command disables EDP on slot 1, ports 2 and 4 on a modular switch:

```
disable edp ports 1:2, 1:4
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 4.0 to support modular switches.

This command was modified in ExtremeWare 6.1 to support the `all` parameter.

Platform Availability

This command is available on all platforms.

disable flooding ports

```
disable flooding ports <portlist>
```

Description

Disables packet flooding on one or more ports.

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

N / A.

Usage Guidelines

Flooding configures the specified ports to act like a hub. Disabling flooding means that only broadcast traffic, EDP traffic, and packets destined to a permanent MAC address matching that port number, are forwarded.

Disabling flooding does not automatically enable learning on the port: use the `enable learning ports` command to re-enable learning on the specified ports.

Learning and flooding are mutually exclusive. To enable learning, you must disable flooding.

On a modular switch, `<portlist>` can be a list of slots and ports. On a stand-alone switch, `<portlist>` can be one or more port numbers. For a detailed explanation of port specification, see “Modular Switch Numerical Ranges” or “Stand-alone Switch Numerical Ranges” in Chapter 1.

Example

The following command disables flooding on ports6, 7 and 8 on a stand-alone switch:

```
disable flooding ports 6,7,8
```

History

This command was first available in ExtremeWare 6.2.1.

Platform Availability

This command is available on *i*-series platforms.

disable g1-module support

```
disable g1-module support
```

Description

Disables I/O modules that do not have the "*i*" series chipset in a BlackDiamond switch.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

To increase the performance of multicast applications in a BlackDiamond switch, you can disable I/O modules that do not have the "*i*" series chipset. When you disable support for older modules (without the "*i*" series chipset), the modules are not powered up, and they do not pass traffic in a BlackDiamond system.

You must save the configuration and reboot the switch for these changes to take effect.

Example

The following command enables I/O modules that do not have the "*i*" series chipset:

```
enable g1-module support
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on BlackDiamond switches.

disable jumbo-frame ports

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

Description

Disables jumbo frame support on a port.

For PoS modules, this command applies to PoS ports when disabling jumbo-frame support changes the negotiated maximum receive unit (MRU) size.

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
all	Specifies all ports on the switch.

Default

Disabled.

Usage Guidelines

On a modular switch, `<portlist>` can be a list of slots and ports. On a stand-alone switch, `<portlist>` can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

Use the `disable jumbo-frame ports` command when you no longer need jumbo frame support.

Example

The following command disables jumbo frame support on port 4 on a stand-alone switch:

```
disable jumbo-frame ports 4
```

The following command disables jumbo frame support on slot 1, port 2 on a BlackDiamond switch:

```
disable jumbo-frame 1:2
```

History

This command was first available in ExtremeWare 6.0.

This command was modified in ExtremeWare 6.1 to support PoS modules.

Platform Availability

This command is available on the "*i*" series platforms.

disable learning ports

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

Description

Disables MAC address learning on one or more ports for security purposes.

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

Enabled.

Usage Guidelines

On a modular switch, `<portlist>` can be a list of slots and ports. On a stand-alone switch, `<portlist>` can be one or more port numbers.For a detailed explanation of port specification, see “Modular Switch Numerical Ranges” or “Stand-alone Switch Numerical Ranges” in Chapter 1.

If MAC address learning is disabled, only broadcast traffic, EDP traffic, and packets destined to a permanent MAC address matching that port number, are forwarded.

Use this command in a secure environment where access is granted via permanent forwarding databases (FDBs) per port.

Learning must be disabled to allow port flooding. See the `enable flooding` command for information on enabling port flooding.

Example

The following command disables MAC address learning on port 4 on a stand-alone switch:

```
disable learning ports 4
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 6.1 to support PoS modules.

Platform Availability

This command is available on all platforms.

disable mirroring

```
disable mirroring
```

Description

Disables port-mirroring.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Use the `disable mirroring` command to stop configured copied traffic associated with one or more ports.

Example

The following command disables port-mirroring:

```
disable mirroring
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

disable ports

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

Description

Disables one or more ports on the switch.

For PoS modules, brings down the PPP link on the specified port and changes the port status LED to blinking green.

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
all	Specifies all ports on the switch.

Default

Enabled.

Usage Guidelines

On a modular switch, `<portlist>` can be a list of slots and ports. On a stand-alone switch, `<portlist>` can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

Use this command for security, administration, and troubleshooting purposes.

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

Example

The following command disables ports 3, 5, and 12 through 15 on a stand-alone switch:

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

The following command disables slot 1, ports 3, 5, and 12 through 15 on a modular switch:

```
disable ports 1:3,1:5,1:12-1:15
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 4.0 to support modular switches.

This command was modified in ExtremeWare 6.1 to support PoS modules.

Platform Availability

This command is available on all platforms.

disable sharing

```
disable sharing [<port>]
```

Description

Disables a load-sharing group of ports.

Syntax Description

port	Specifies the master port of a load-sharing group. On a modular switch, is a combination of the slot and port number, in the format <slot>:<port>.

Default

Disabled.

Usage Guidelines

This command increases bandwidth tracking and resiliency.

On a modular switch, <port> is specified as <slot>:<port number>. On a stand-alone switch, <port> is the port configured as the load-sharing master port. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

Example

The following command disables sharing on master logical port 9, which contains ports 9-12 on a stand-alone switch:

```
disable sharing 9
```

The following command disables sharing on master logical port 9 in slot 3, which contains ports 9 through 12 on a modular switch:

```
disable sharing 3:9
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 4.0 to support modular switches.

Platform Availability

This command is available on all platforms.

disable slot

```
disable slot [<slot number> | all]
```

Description

Disables one or all slots on a BlackDiamond or Alpine switch, and leaves the blade in a power down state.

Syntax Description

slot number	Specifies the slot to be disabled.
all	Species that all slots in the device should be disabled.

Default

Enabled.

Usage Guidelines

This command allows the user to disable a slot. When the user types this command, the I/O card in that particular slot number is brought down, and the slot is powered down. The LEDs on the card go OFF.

A disabled slot can be re-enabled using the `enable slot` command.

The `show slot` command, if invoked after the user disables the slot, shows this slot state as “Disabled.” The user can either disable a slot individually or use the `disable slot all` to disable all the slots.

If there is no I/O card present in a slot when the user disables the slot, the slot still goes to the “Disable” state. If a card is inserted in a slot that has been disabled, the card does not come up and stays in the “disabled” state until the slot is enabled by using the `enable slot` command. below.

If you do not save the configuration before you do a switch reboot, the slot will be re-enabled upon reboot. If you save the configuration after disabling a slot, the slot will remain disabled after a reboot.

Example

The following command disables slot 5 on the switch:

```
disable slot 5
```

History

This command was first available in ExtremeWare 6.2.1.

Platform Availability

This command is available on BlackDiamond and Alpine switches only.

disable smartredundancy

```
disable smartredundancy [<portlist>]
```

Description

Disables the smart redundancy feature.

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

Disabled.

Usage Guidelines

On a modular switch, <portlist> can be a list of slots and ports. On a stand-alone switch, <portlist> can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

Use with Extreme Networks switches that support privacy and backup uplinks.

When smartredundancy is disabled, the switch changes the active link only when the current active link becomes inoperable.

Example

The following command disables the smart redundancy feature on ports 1-4:

```
disable smartredundancy 1-4
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms switches.

enable edp ports

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

Description

Enables the Extreme Discovery Protocol (EDP) on one or more ports.

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
all	Specifies all ports on the switch.

Default

Enabled.

Usage Guidelines

On a modular switch, `<portlist>` can be a list of slots and ports. On a stand-alone switch, `<portlist>` can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

EDP is useful when Extreme Networks switches are attached to a port.

The EDP is used to locate neighbor Extreme Networks switches and exchange information about switch configuration. When running on a normal switch port, EDP is used to by the switches to exchange topology information with each other. Information communicated using EDP includes the following:

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

For version 2.0 and 4.0:

Information communicated using EDP also includes the following:

- Virtual chassis identifier and port number
- Listing of all virtual chassis identifiers

Example

The following command enables EDP on port 7 on a stand-alone switch:

```
enable edp ports 7
```

The following command enables EDP on slot 1, port 3 on a modular switch:

```
enable edp ports 1:3
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 4.0 to support modular switches.

Platform Availability

This command is available on all platforms.

enable flooding ports

```
enable flooding ports <portlist>
```

Description

Enables packet flooding on one or more ports.

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

Ports are enabled for learning, not flooding.

Usage Guidelines

This command configures the specified ports to act like a hub. When flooding is enabled on a particular port, *all* frames and packets are passed on to other member ports that have flooding enabled. This includes all broadcast, multicast, known unicast and unknown unicast packets (including EPD). To make effective use of this feature you should have flooding enabled on more than one port.

Learning and flooding are mutually exclusive. To enable flooding, you must first disable learning.

When ports are configured for flooding, the FDB will be flushed for the entire system, which means all the entries in the dynamic FDB must be relearned.

On a modular switch, `<portlist>` can be a list of slots and ports. On a stand-alone switch, `<portlist>` can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

Example

The following command enables flooding on ports6, 7 and 8 on a stand-alone switch:

```
enable flooding ports 6,7,8
```

History

This command was first available in ExtremeWare 6.2.1.

Platform Availability

This command is available on *i*-series platforms.

enable g1-module support

```
enable g1-module support
```

Description

Enables I/O modules that do not have the *"i"* series chipset in a BlackDiamond switch.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

You must save the configuration and reboot the switch for a changes to take effect.

Example

The following command enables I/O modules that do not have the *"i"* series chipset:

```
enable g1-module support
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on BlackDiamond switches.

enable jumbo-frame ports

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

Description

Enables support on the physical ports that will carry jumbo frames.

For PoS modules, enables jumbo-frame support to specific PoS ports when jumbo-frame support changes the negotiated maximum receive unit (MRU) size.

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
all	Specifies all ports on the switch.

Default

Disabled.

Usage Guidelines

Increases performance to back-end servers or allows for VMAN 802.1q encapsulations.

You must configure the maximum MTU size of a jumbo frame before you can use the `enable jumbo-frame ports` command. Use the `config jumbo-frame size` command to configure the MTU size.

On a modular switch, `<portlist>` can be a list of slots and ports. On a stand-alone switch, `<portlist>` can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

Example

The following command enables jumbo frame support on port 5 on a stand-alone switch:

```
enable jumbo-frame ports 5
```

The following command enables jumbo frame support on slot 3, port 5 on a modular switch:

```
enable jumbo-frame ports 3:5
```

History

This command was first available in ExtremeWare 6.0.

This command was modified in ExtremeWare 6.1 to support PoS modules.

Platform Availability

This command is available on the "*i*" series platforms.

enable learning ports

```
enable learning ports <portlist>
```

Description

Enables MAC address learning on one or more ports.

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

Enabled.

Usage Guidelines

On a modular switch, <portlist> can be a list of slots and ports. On a stand-alone switch, <portlist> can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

Example

The following command enables MAC address learning on ports 7 and 8 on a stand-alone switch:

```
enable learning ports 7,8
```

The following command enables MAC address learning on slot 1, ports 7 and 8 on a modular switch:

```
enable learning ports 1:7-8
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 4.0 to support modular switches.

Platform Availability

This command is available on all platforms.

enable mirroring

```
enable mirroring to port [<portlist>] [tagged | untagged]
```

Description

Dedicates a port on the switch to be the mirror output port.

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
tagged	Configures the ports as tagged.
untagged	Configures the ports as untagged.

Default

N/A.

Usage Guidelines

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

- Physical port — All data that traverses the port, regardless of VLAN configuration, is copied to the monitor port.
- VLAN — All data to and from a particular VLAN, regardless of the physical port configuration, is copied to the monitor port.
- Virtual port — All data specific to a VLAN on a specific port is copied to the monitor port.

Up to eight mirroring filters and one monitor port can be configured on the switch. After a port has been specified as a monitor port, it cannot be used for any other function. Frames that contain errors are not mirrored.

For version 6.0 and higher:

- `tagged` and `untagged` are added to the command syntax.

For version 4.0 and higher:

- `to` is added to the command syntax.
- Supports modular switches.

On a modular switch, `<portlist>` can be a list of slots and ports. On a stand-alone switch, `<portlist>` can be one or more port numbers. For a detailed explanation of port specification, see “Modular Switch Numerical Ranges” or “Stand-alone Switch Numerical Ranges” in Chapter 1.

For version 2.0 and 4.0:

- In addition to the physical port, VLAN, and virtual port, the traffic filter can be defined based on the following criteria:

— MAC source address / destination address — All data sent to or received from a particular source or destination MAC address is copied to the monitor port.

For MAC mirroring to work correctly, the MAC address must already be present in the forwarding database (FDB).

Example

The following example selects port 3 as a tagged mirror port on a stand-alone switch:

```
enable mirroring to ports 3 tagged
```

The following example selects slot 1, port 3 as the mirror port on a modular switch:

```
enable mirroring to ports 1:3
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 4.0 to support the `tag | untagged` keywords and modular switches.

Platform Availability

This command is available on all platforms.

enable ports

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

Description

Enables a port.

For PoS modules, enables the PPP link on the specified port, and changes the port status LED to solid green (if no other problems exist).

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
all	Specifies all ports on the switch.

Default

All ports are enabled.

Usage Guidelines

Use this command to enable the port(s) if you disabled the port(s) for security, administration, or troubleshooting purposes.

On a modular switch, `<portlist>` can be a list of slots and ports. On a stand-alone switch, `<portlist>` can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

Example

The following command enables ports 3, 5, and 12 through 15 on the stand-alone switch:

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

The following command enables slot 1, ports 3, 5, and 12 through 15 on the modular switch:

```
enable ports 1:3, 1:5, 1:12-1:15
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 6.1 to support PoS modules.

This command was modified in ExtremeWare 4.0 to support the modular switches.

Platform Availability

This command is available on all platforms.

enable sharing grouping

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

Description

Defines a load-sharing group of ports. The ports specified in <portlist> are grouped to the master port.

Syntax Description

port number	Specifies the master port for a loadsharing group.
portlist	Specifies one or more ports or slots and ports to be grouped to the master port. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
port-based	Uses the ingress port as criteria for egress port selection. (6.0 and higher)
address-based	Uses addressing information as criteria for egress port selection. (6.0 and higher)
round-robin	Forwards packets to all egress ports in a round-robin fashion. (6.0 and higher)

Default

Disabled.

Usage Guidelines

On a modular switch, <portlist> can be a list of slots and ports. On a stand-alone switch, <portlist> can be one or more port numbers. For a detailed explanation of port specification, see “Modular Switch Numerical Ranges” or “Stand-alone Switch Numerical Ranges” in Chapter 1.

Load sharing allows you to increase bandwidth and resilience between switches 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 or a “master” port. For example, VLANs see the load-sharing group as a single logical port. The algorithm also guarantees 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.

Load sharing must be enabled on both ends of the link, or a network loop will result.

Load sharing is most useful in cases where the traffic transmitted from the switch to the load-sharing group is sourced from an equal or greater number of ports on the switch. For example, traffic transmitted to a two-port load-sharing group should originate from a minimum of two other ports on the same switch.

This feature is supported between Extreme Networks switches only, but may be compatible with third-party “trunking” or sharing algorithms. Check with an Extreme Networks technical representative for more information.

Modular switch load-sharing groups are defined according to the following rules:

- Ports on the I/O module are divided into groups of two or four.
- Port in a load-sharing group must be part of the same I/O module.
- Ports in a load-sharing group must be contiguous.

Summit switch load-sharing groups are defined according to the following rules:

- Ports in a load-sharing group must be contiguous.
- Ports on the switch are divided into groups of two or four.
- Address-based and round-robin load sharing algorithms do not apply.

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 ports assigned to a load-sharing group cannot be members of a Spanning Tree Domain (STPD).

When using load sharing, you should always reference the master logical port of the load-sharing group 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.

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 receive a valid indication that the port is not in a forwarding state, and the partner switch will continue to forward packets.

For version 6.0 and higher:

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 one of three load-sharing algorithms on the switch, as follows:

- Port-based — Uses the ingress port to determine which physical port in the load-sharing group is used to forward traffic out of the switch.
- Address-based — 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, and the TCP port number.
  - IPX packets — Uses the source and destination MAC address, and IPX network identifiers.
  - All other packets — Uses the source and destination MAC address.
- Round-robin — When the switch receives a stream of packets, it forwards one packet out of each physical port in the load-sharing group using a round-robin scheme.

Using the round-robin algorithm, packet sequencing between clients is not guaranteed.

If you do not explicitly select an algorithm, the port-based scheme is used. However, the address-based algorithm has a more even distribution and is the recommended choice.

The address-based and round-robin load-sharing algorithms are supported by BlackDiamond switch modules that use the "*i*" chipset and all Alpine 3800 switch modules. The modules end with an "i" in their model designation (for example, G12SXi), and require the use of the MSM64i. The address-based and round-robin load-sharing algorithms are also supported by Summit switches that use the "*i*" chipset, such as the Summit 7i switch.

For BlackDiamond and Alpine switches (version 6.0 and higher):

To set up a modular switch to load share among ports, you must create a load-sharing group of ports. Load-sharing groups are defined according to the following rules:

- 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.

On I/O modules that do not have the “*i*” chipset, the following additional rules apply:

- Ports in a load-sharing group must be contiguous.
- Ports on the I/O module are divided into groups of two or four.

For Summit switches (version 6.0 and higher):

On switches that do not have the “*i*” chipset, the following additional rules apply:

- Ports in a load-sharing group must be contiguous.
- Ports on the switch are divided into groups of two or four.
- Address-based and round-robin load sharing algorithms do not apply.

On Summit switches other than Summit1, Summit2, Summit3, Summit4/FX, Summit24, and Summit48, the following rules apply:

- A group can contain up to 8 ports.
- The ports in a group do not need to be contiguous.

Example

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

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

The following example defines a load-sharing group on slot 3 that contains ports 9 through 12, and uses the first port in the group as the master logical port 9 on a modular switch:

```
enable sharing 3:9 grouping 3:9-3:12
```

In this example, logical port 3:9 represents physical ports 3:9 through 3:12.

History

This command was first available in ExtremeWare 2.0.

The command was modified in ExtremeWare 4.0 to support modular switches.

The command was modified in ExtremeWare 6.0 to support the `algorithm` parameter.

Platform Availability

This command is available on all platforms.

enable slot

```
enable slot [<slot number> | all]
```

Description

Enables one or all slots on a BlackDiamond or Alpine switch.

Syntax Description

slot number	Specifies the slot to be enabled.
all	Species that all slots in the device should be enabled.

Default

Enabled.

Usage Guidelines

This command allows the user to enable a slot that has been previously disabled using the `disable slot` command.

When the user enters the enable command, the disabled I/O card in the specified slot is brought up, and the slot is made operational, if possible, or goes to the appropriate state as determined by the card state machine. The LEDs on the card are brought ON as usual. The user can either enable a slot individually, or use the `enable slot all` command to enable all the slots.

After the user enables the slot, the `show slot` command shows the state as “Operational” or will display the appropriate state if the card could not be brought up successfully. Note that there is no card state named “Enable” and the card goes to the appropriate states as determined by the card state machine when the `enable slot` command is invoked.

Only slots that have their state as “disabled” can be enabled using this command. If this command is used on slots that are in states other than “disabled,” the card state machine takes no action on these slots.

Example

The following command enables slot 5 on the switch:

```
enable slot 5
```

History

This command was first available in ExtremeWare 6.2.1.

Platform Availability

This command is available on BlackDiamond and Alpine switches only.

enable smartredundancy

```
enable smartredundancy <portlist>
```

Description

Enables the Smart Redundancy feature on the redundant Gigabit Ethernet port.

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

Enabled.

Usage Guidelines

When the Smart Redundancy feature is enabled, the switch always uses the primary link when the primary link is available.

On a modular switch, <portlist> can be a list of slots and ports. On a stand-alone switch, <portlist> can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

Example

The following command enables the Smart Redundancy feature on port 4 on a switch:

```
enable smartredundancy 4
```

The following command enables the Smart Redundancy feature on slot 1, port 4 on a BlackDiamond switch:

```
enable smartredundancy 1:4
```

History

This command was first available in ExtremeWare 2.0.

Support for modular switches was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

restart ports

```
restart ports [<portlist>
```

Description

Resets autonegotiation for one or more ports by resetting the physical link.

For PoS modules, causes the PPP link to be renegotiated.

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

N / A.

Usage Guidelines

On a modular switch, <portlist> can be a list of slots and ports. On a stand-alone switch, <portlist> can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

Example

The following command resets autonegotiation on port 4 on a stand-alone switch:

```
restart ports 4
```

The following command resets autonegotiation on slot 1, port 4 on a modular switch:

```
restart ports 1:4
```

History

This command was first available in ExtremeWare 4.0.

This command was modified in ExtremeWare 6.1 to support PoS modules.

Platform Availability

This command is available on all platforms.

show edp

```
show edp {<portlist>}
```

Description

Displays connectivity and configuration information for neighboring Extreme Networks switches.

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

N / A.

Usage Guidelines

On a modular switch, <portlist> can be a list of slots and ports. On a stand-alone switch, <portlist> can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

Use the show edp command to display neighboring switches and configurations. This is most effective with Extreme Networks switches.

Example

The following command displays the connectivity and configuration of neighboring Extreme Networks switches:

```
show edp
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 4.0 to support modular switches.

Platform Availability

This command is available on all platforms.

show mirroring

```
show mirroring
```

Description

Displays the port-mirroring configuration on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

You must configure mirroring on the switch to display mirroring statistics. Use the `show mirroring` command to configure mirroring.

You can use this command to display mirroring statistics and determine if mirroring is enabled or disabled on the switch.

To view the status of port-mirroring on the switch, use the `show mirroring` command. The `show mirroring` command displays information about the enable/disable state for port-mirroring.

Example

The following command displays switch mirroring statistics:

```
show mirroring
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

show ports collisions

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

Description

Displays real-time collision statistics.

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

N / A

Usage Guidelines

If you do not specify a port number or range of ports, collision statistics are displayed for all ports.

On a modular switch, <portlist> can be a list of slots and ports. On a stand-alone switch, <portlist> can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

This status information may be useful for your technical support representative if you have a network problem.

Example

The following command displays real-time collision statistics on port 7 on a stand-alone switch:

```
show ports 7 collisions
```

The following command displays real-time collision statistics on slot 1, port 7 on a modular switch:

```
show ports 1:7 collisions
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 4.0 to support modular switches.

Platform Availability

This command is available on all platforms.

show ports configuration

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

Description

Displays port configuration statistics.

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

N/A

Usage Guidelines

On a modular switch, <portlist> can be a list of slots and ports. On a stand-alone switch, <portlist> can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

If you do not specify a port number or range of ports, configuration statistics are displayed for all ports.

This status information may be useful for your technical support representative if you have a network problem.

This command displays port configuration, which includes:

- Port state
- Link state
- Link speed
- Duplex mode
- Flow control
- Load sharing information
- Link media information

For version 6.0 and higher:

- Auto on/off

Example

The following command displays the port configuration statistics for all ports on a switch:

```
show ports config
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 4.0 to support modular switches.

Platform Availability

This command is available on all platforms.

show ports info

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

Description

Displays detailed system-related information.

For PoS modules, displays port information that includes new DiffServ and RED configuration parameters.

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
detail	Specifies detailed port information. (6.0 and higher)

Default

N/A.

Usage Guidelines

If you do not specify a port number or range of ports, detailed system-related information is displayed for all ports. The data is displayed in a table format.

On a modular switch, `<portlist>` can be a list of slots and ports. On a stand-alone switch, `<portlist>` can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

This status information may be useful for your technical support representative if you have a network problem.

For version 6.0 and higher:

- The `detail` parameter is used to provided more specific port information. The data is called out with written explanations versus displayed in a table format.

Example

The following command displays port system-related information:

```
show ports info
```

Following is the output from this command:

```
Port         Diag   Flags     Link    Link  Num   Num   Num   Jumbo   QOS    Load
                              State   UPs   STP   VLAN  Proto Size
  29           P  e--m-------E ready   0     1     1     1    9216
  30           P  e--m-------E ready   0     1     1     1    9216
  31           P  e--m-------E ready   0     1     1     1    9216
  32           P  e--m-------E active  1     1     1     1    9216
  33           P  e--m-------E ready   0     1     1     1    9216
  34           P  e--m-------D ready   0     1     1     1    9216
```

```
  35          P  e--m-------D ready   0     1     1    1    9216
  36          P  e--m-------D active  1     1     1    1    9216
  37          P  e--m-------D ready   0     1     1    1    9216
  38          P  e--m-------D ready   0     1     1    1    9216
  39          P  e--m-------D ready   0     1     1    1    9216
  40          P  e--m-------D ready   0     1     1    1    9216
  41          P  e--m-------D ready   0     1     1    1    9216
  42          P  e--m-------E ready   0     1     1    1    9216
  43          P  e--m-------E ready   0     1     1    1    9216
  44          P  e--m-------E ready   0     1     1    1    9216
  45          P  e--m-------E ready   0     1     1    1    9216
  46          P  e--m-------E active  1     1     1    1    9216
  47          P  e--m-------E ready   0     1     1    1    9216
  48          P  e--m-------E ready   0     1     1    1    9216
  49          P  e--m-------E ready   0     1     1    1    9216
  50          P  e--m-------E ready   0     1     1    1    9216

Flags : d - DLCS Enabled, e - Extreme Discovery Protocol Enabled, g - Egress TOS
        Enabled, j - Jumbo Frame Enabled, l - Load Sharing Enabled, m - MAC
        Learning Enabled, n - Ingress TOS Enabled, o - Dot1p Vlan Priority
        Replacement Enabled, q - Background QOS Monitoring Enabled.
        a - Load Sharing Algorithm address-based, p - Load Sharing Algorithm
        port-based, r - Load Sharing Algorithm round-robin, G - SLB GoGo Mode
        h - Hardware Redundant Phy, P - Software Primary Port, R - Software
        Redundant Port, f - Flooding Enabled, D - Port Disabled, E - Port Enabled
Diag : P - Passed, F - Failed
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 4.0 to support modular switches.

This command was modified in ExtremeWare 6.0 to support the `detail` keyword.

This command was modified in ExtremeWare 6.1 to support PoS modules.

This command was modified in ExtremeWare 6.2.2 to indicate disabled or enabled status.

Platform Availability

This command is available on all platforms.

show ports sharing

```
show ports sharing
```

Description

Displays port loadsharing groups.

Syntax Description

This command has no arguments or variables..

Default

N/A

Usage Guidelines

Example

The following command displays the port loadsharing groups:

```
show ports sharing
```

```
MSM64: 7 # show ports sharing
Load Sharing Monitor
Config    Current      Ld Share     Ld Share  Link      Link
Master    Master       Algorithm    Group     Status    Ups
=================================================================
4:1                       p            4:1      R        0
                          p            4:2      R        0
                          p            4:3      R        0

Flags: Link Status: A-Active, R-Ready, D-Disabled, NP-Not Present
        Ld Sh Algo: p-port based, a-address based, r-round robin
```

History

This command was first available in ExtremeWare 6.2.2.

Platform Availability

This command is available on all platforms.

show ports packet

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

Description

Displays a histogram of packet statistics.

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

N/A.

Usage Guidelines

If you do not specify a port number or range of ports, a histogram is displayed for all ports.

On a modular switch, <portlist> can be a list of slots and ports. On a stand-alone switch, <portlist> can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

This status information may be useful for your technical support representative if you have a network problem.

The following packet statistics are displayed:

- port number
- link status
- packet size

Example

The following command displays packet statistics for ports 1 through 3 on a stand-alone switch:

```
show ports 1-3 packet
```

The following command displays packet statistics for slot 1, ports 1 through 3 on a modular switch:

```
show ports 1:1-1:3 packet
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 4.0 to support modular switches.

Platform Availability

This command is available on all platforms.

show ports utilization

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

Description

Displays real-time port utilization information.

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

N / A.

Usage Guidelines

Use the [Spacebar] to toggle between packet, byte, and bandwidth utilization information.

On a modular switch, <portlist> can be a list of slots and ports. On a stand-alone switch, <portlist> can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

If you do not specify a port number or range of ports, port utilization information is displayed for all ports.

This status information may be useful for your technical support representative if you have a network problem.

Example

The following command displays utilization statistics for port 1 on a stand-alone switch:

```
show ports 1 utilization
```

The following command displays utilization statistics for slot 3, port 1 on a modular switch:

```
show ports 3:1 utilization
```

The following examples show the output from the show ports utilization command for all ports on the switch. The three displays show the information presented when you use the spacebar to toggle through the display types. The first display shows utilization in terms of packets:

```
Link Utilization Averages                          Wed Jan 23 21:29:45 2002
Port         Link    Receive      Peak Rx       Transmit       Peak Transmit
             Status  packet/sec    pkt/sec       pkt/sec        pkt/sec
================================================================================
   1           A       43         255             4             14
   2           R        0           0             0              0
   3           R        0           0             0              0
   4           R        0           0             0              0
   5           R        0           0             0              0
   6           R        0           0             0              0
   7           R        0           0             0              0
   8           R        0           0             0              0

================================================================================
           Link Status: A-Active R-Ready D-Disabled NP-Not Present
        spacebar->toggle screen U->page up  D->page down ESC->exit
```

The second display shows utilization in terms of bytes:

```
Link Utilization Averages                          Wed Jan 23 21:30:03 2002
Port    Link    Receive      Peak Rx       Transmit       Peak Transmit
        Status  bytes/sec     bytes/sec     bytes/sec      bytes/sec
================================================================================
   1      A      1102       69555            536           2671
   2      R         0           0              0              0
   3      R         0           0              0              0
   4      R         0           0              0              0
   5      R         0           0              0              0
   6      R         0           0              0              0
   7      R         0           0              0              0
   8      R         0           0              0              0

================================================================================
           Link Status: A-Active R-Ready D-Disabled NP-Not Present
```

The third display shows bandwidth utilization:

```
Link Utilization Averages                          Wed Jan 23 21:30:19 2002
Port     Link     Link   Receive      Peak Rx     Transmit     Peak Transmit
         Status   Speed  % bandwidth  % bandwidth  % bandwidth  % bandwidth
================================================================================
   1      A       100    0.00          0.60         0.00         0.02
   2      R              0.00          0.00         0.00         0.00
   3      R              0.00          0.00         0.00         0.00
   4      R              0.00          0.00         0.00         0.00
   5      R              0.00          0.00         0.00         0.00
   6      R              0.00          0.00         0.00         0.00
   7      R              0.00          0.00         0.00         0.00
   8      R              0.00          0.00         0.00         0.00


================================================================================
           Link Status: A-Active R-Ready D-Disabled NP-Not Present
        spacebar->toggle screen U->page up  D->page down ESC->exit
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 4.0 to support modular switches.

Platform Availability

This command is available on all platforms.

show sharing address-based

```
show sharing address-based
```

Description

Displays the address-based load sharing configuration.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This feature is available using the address-based load-sharing algorithm only. 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 address, and the TCP port number.
- IPX packets—Uses the source and destination MAC address and IPX identifiers.
- All other packets—Uses the source and destination MAC address.

To verify your configuration, use the `show sharing address-based` command. The show sharing address-based output displays the addressed-based configurations on the switch.

Example

The following example displays the address-based load sharing configuration on the switch:

```
show sharing address-based
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platform and the Alpine 3800 series switch modules.

show slot

```
show slot <slot number>
```

Description

Displays the slot-specific information.

For PoS and MPLS modules, displays information that includes data about the software images loaded on the module, as well as status information on the network processors and GPP.

Syntax Description

slot number	Specifies a slot on a modular switch.

Default

N/A.

Usage Guidelines

The `show slot` command displays the following information:

- The name of the module installed in the slot
- The serial number of the module
- The part number of the module
- The state of the module, whether the power is down, if the module is operational, if a diagnostic being run, if there is a mismatch between the slot configuration and the module in the slot
- The status of the ports on the module

If you do not specify a slot number, information for all slots is displayed.

For PoS and MPLS modules:

The ExtremeWare technology release that supports the PoS or MPLS module includes multiple software packages. One software package runs on the MSM module while another package runs on each PoS or MPLS module. You must download the software packages independently using the ExtremeWare `download image` command. Each software package has an associated version number that you can display using the `show version` command. It is recommended (not required), that the MSM software package and the PoS or MPLS module software package be the same version. To ensure compatibility, the MSM performs an automatic compatibility check before a PoS or MPLS module is activated. If the versions of the software packages are incompatible, the PoS or MPLS ports on the module will not come up and the `show slot` command will indicate that the software on the PoS or MPLS module is incompatible with the MSM software.

Assuming the PoS or MPLS module has no problems, the command `show slot <slot>` (where *"<slot>"* is the number of the slot where you installed the module) displays that ExtremeWare has detected the module and set it to the OPERATIONAL state.

As the module progresses through its initialization, the `show slot <slot>` command displays the GPP subsystem change state to OPERATIONAL, and then each of the network processors will change state to OPERATIONAL.

When the GPP subsystem completes its initialization cycle and the subsystem state is OPERATIONAL, use the `show diagnostics {<slot>}` *command to check the results of the module power-on self test (POST).*

If the STATUS LED on the PoS MPLS module turns amber and blinks, use the `show slot <slot>` command to display the slot status information. The `show slot <slot>` command also displays operational information related to the PoS MPLS module. Information displayed includes the BlackDiamond switch fabric card state, Network Processor status, General Purpose Processor status, hardware serial number and type, and image version and boot settings.

For the PoS and MPLS modules, the information displayed by this command includes data about the software images loaded on the module and information about the operational status and backplane connections of the module.

Example

The following example displays I/O module information for an I/O module in slot 4:

```
show slot 4
```

History

This command was first available in ExtremeWare 4.0.

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12 to support MPLS modules.

This command was modified in ExtremeWare 6.1 to support PoS modules.

Platform Availability

This command is available on modular switches only.

unconfig ports display string

```
unconfig ports <portlist> display-string
```

Description

Clears the user-defined display string from one or more ports.

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

N/A.

Usage Guidelines

This command removes the display string that you configured using the `config ports display-string` command.

On a modular switch, `<portlist>` can be a list of slots and ports. On a stand-alone switch, `<portlist>` can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

Example

The following command clears the user-defined display from port 4 on a stand-alone switch:

```
unconfig ports 4 display-string
```

The following command clears the user-defined display string from slot 2, port 4 on a modular switch:

```
unconfig ports 2:4 display-string
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

unconfig ports redundant

```
unconfig ports [<portlist> | <port id> | mgmt] redundant
```

Description

Clears a previously configured software-controlled redundant port.

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
portid	Specifies a port using the display string configured for the port. Only one port can be specified using this method.
mgmt	Specifies the management port. Supported only for switches that provide a management port.

Default

N / A.

Usage Guidelines

On a modular switch, `<portlist>` can be a list of slots and ports. On a stand-alone switch, `<portlist>` can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

The `<port id>` is the display string configured for the port. Use the `config ports <portnumber> display-string <string>` command to configure a display string for the port.

The list of port numbers or the port display string specifies the redundant port(s).

Example

The following command unconfigures a software-controlled redundant port on a stand-alone switch:

```
unconfig ports 4 redundant
```

The following command unconfigures a software-controlled redundant port on a modular switch:

```
unconfig ports 2:3 redundant
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

unconfig slot

```
unconfig slot <slot number>
```

Description

Clears a slot of a previously assigned module type.

Syntax Description

slot number	Specifies a slot on a modular switch.

Default

N/A.

Usage Guidelines

For PoS and MPLS modules, clears a slot of a previously assigned PoS or MPLS module and removes any port-related information associated with the slot.

In the ExtremeWare IP Technology Services Release based on v6.1.8b12:

- Keywords were added to specify the PoS modules (`p12c, p3c`) and the MPLS module (`mpls`).

Example

The following command clears slot 4 of a previously assigned module type:

```
unconfig slots 4
```

The following command clears slot 3 of a previously assigned MPLS module and removes any port-related information associated with the slot:

```
unconfig slot 3
```

History

This command was first available in ExtremeWare 4.0.

This command was available in an ExtremeWare IP Technology Services Release based on v6.1.8b12 to support MPLS modules.

This command was modified in an ExtremeWare IP Technology Services Release based on v6.1.5b20 to support PoS modules.

Platform Availability

This command is available on modular switches only.

5 VLAN Commands

This chapter describes the following commands:

- Commands for creating and deleting VLANs and performing basic VLAN configuration
- Commands for defining protocol filters for use with VLANs
- Commands for enabling or disabling the use of Generic VLAN Registration Protocol (GVRP) information on a switch and its ports

VLANs can be created according to the following criteria:

- Physical port – A port-based VLAN consists of a group of one or more ports on the switch. A port can be a member of only one port-based VLAN, and is by default a member of the VLAN named "Default."
- 802.1Q tag – Tagging is most commonly used to create VLANs that span switches.
- Ethernet, LLC SAP, or LLC/SNAP Ethernet protocol type – Protocol-based VLANs are most often used in situations where network segments contain hosts running multiple protocols.
- A combination of these criteria

The Generic VLAN Registration Protocol (GVRP) allows switches to learn some VLAN information automatically instead of requiring manual configuration in each switch. A VLAN can provide GVRP information about its VLANs and accept information about VLANs from other GVRP-enabled switches. Depending on the circumstances, information learned in this manner may cause ports to be added to VLANs already existing on the switch, or may cause new tagged VLANs to be created automatically.

NOTE

GVRP is not supported in ExtremeWare versions 6.1 or later.

config cpu-transmit-priority

```
config cpu-transmit-priority [high | normal]
```

Description

Configures the CPU transmit priority.

Syntax Description

high	Sets the CPU transmit priority to high.
normal	Sets the CPU transmit priority to normal.

Default

High.

Usage Guidelines

The CPU transmit priority is set to `high` by default to control the priority in which packets are transmitted from the switch in the event that lower priority queues are congested. This mechanism uses internal resources and limits the number of VLANs that can be configured on the switch.

To configure more than 1024 VLANs, set the `cpu-transmit-priority` level to `normal`. The maximum number of VLANs supported is 3000.

Example

The following command, followed by a switch reboot, changes the CPU transmit priority to normal:

```
config cpu-transit-priority normal
```

The following command, followed by a switch reboot, returns the CPU transmit priority to high:

```
config cpu-transit-priority high
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all platforms.

config dot1q ethertype

```
config dot1q ethertype <ethertype>
```

Description

Configures an IEEE 802.1Q Ethertype.

Syntax Description

ethertype	Specifies an Ethertype value.

Default

Ethertype value of 8100.

Usage Guidelines

Use this command if you need to communicate with a switch that supports 802.1Q, but uses an Ethertype value other than 8100. This feature is useful for VMAN tunneling.

Extreme switches assume an Ethertype value of 8100.

You must reboot the switch for this command to take effect.

Example

The following command, followed by a switch reboot, changes the Ethertype value to 9100:

```
config dot1q ethertype 9100
```

History

This command was first available in ExtremeWare 1.0.

Platform Availability

This command is available on all platforms.

config gvrp

```
config gvrp {listen | send | both | none} port <portlist>
```

Description

Configures the sending and receiving of Generic VLAN Registration Protocol (GVRP) information on a port.

Syntax Description

listen	Enables the receipt of GVRP packets on the specified port(s).
send	Enables sending of GVRP packets on the specified port(s).
both	Enables both sending and receiving of GVRP packets.
none	Disables the port from participating in GVRP operation.
portlist	Specifies one or more ports. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

Both sending and receiving.

Usage Guidelines

GVRP must be enabled on the switch as a whole before GVRP data can be sent or received on individual ports.

If GVRP is enabled, send causes information (GVRP packets) about tagged VLANS on the switch to be sent on the specified ports, to neighboring GVRP-enabled switches.

If GVRP is enabled, `listen` means that the switch will receive and act on GVRP information it receives on the specified ports, from neighboring GVRP-enabled switches.

Example

The following commands configure port 3 to receive GVRP information only (by default it can send and listen) and then enables GVRP:

```
config gvrp listen port 3
enable gvrp
```

If the switch receives GVRP information on this port, it will do one of the following:

- If a tagged VLAN already exists with a VLANid that matches the VLANid in the GVRP data, and port 3 is not already a member of that VLAN, add it as a tagged port.
- If no VLAN exists with a VLANid that matches the VLANid in the GVRP data, create a VLAN with the VLANid specified in the GVRP data, and add port 3 as a tagged member port.

History

This command was first available in ExtremeWare 2.0.

This command exists but is not supported in ExtremeWare version 6.1 and later.

Platform Availability

This command is available on all platforms.

config mac-vlan add mac-address

```
config mac-vlan add mac-address [any | <mac_address>] mac-group [any |
<group_number>] vlan <name>
```

Description

Adds a MAC address as a potential member of a MAC-based VLAN.

Syntax Description

mac_address	The MAC address to be added to the specified VLAN. Specified in the form nn:nn:nn:nn:nn:nn.
	any indicates that any MAC-address associated with the specified MAC group may be a member.
group_number	The group number that should be associated with the specified MAC address. Specified as an integer
	any indicates that this MAC address can be associated with any MAC group.
name	The name of the VLAN with which this MAC address should associated.

Default

N/A.

Usage Guidelines

The specified MAC address must be associated with an end station/host only, not a layer-2 repeater device.

Adding a MAC address means that when the specified address is detected on a member port, as specified by its group membership, it can participate in the VLAN.

At least one port must be enabled to use the MAC-based VLAN algorithm before any MAC addresses can be added.

Example

Given ports enabled for MAC-based VLANs as follows:

```
enable mac-vlan mac-group any ports 16,17
enable mac-vlan mac-group 10 ports 11,12
```

The following command sets up the end-station with MAC address 00:00:00:00:00:01 to participate in VLAN engineering via the MAC-enabled ports 16 or 17:

```
config mac-vlan add mac-address 00:00:00:00:00:01 mac-group any vlan engineering
```

MAC address 00:00:00:00:00:01 cannot get access via ports 11 or 12 because it is not configured for mac-group 10.

The following command sets up the endstation 00:00:00:00:00:02 to participate in VLAN engineering through the ports in group 10 (ports 11 or 12) or through ports 16 or 17 (enabled for any mac-group):

```
config mac-vlan add mac-address 00:00:00:00:00:02 mac-group 10 vlan engineering
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

config mac-vlan delete

```
config mac-vlan delete [all | mac-address [<mac_address> | any]]
```

Description

Removes a MAC address from any MAC-based VLANs with which it was associated.

Syntax Description

all	Indicates that all MAC addresses should be removed from all VLANs.
mac_address	The MAC address to be removed. Specified in the form nn:nn:nn:nn:nn:nn. any indicates that all MAC-addresses should be removed from all VLANs.

Default

NA.

Usage Guidelines

None.

Example

The following command removes the endstation with MAC address 00:00:00:00:00:02 from participating in any MAC-based VLANs.

```
config mac-vlan delete mac-address 00:00:00:00:00:02
```

The following commands remove the all MAC addresses from participating in any VLANs:

```
config mac-vlan delete all
config mac-vlan delete mac-address any
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

config ports monitor vlan

```
config ports <portlist> monitor vlan <name >
```

Description

Configures VLAN statistic monitoring on a per-port basis.

Syntax Description

portlist	Specifies one or more ports. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

None.

Example

The following command configures per port monitoring for a set of ports on slot 8 for the VLAN named *accounting*:

```
config ports 8:1-8:6 monitor vlan accounting
```

You can monitor up to four VLANs on the same port by issuing the command four times. For example, if you want to monitor VLANs dog1, dog2, dog3, and dog4 on slot 1, use the following commands:

```
config ports 1:* monitor vlan dog1
config ports 1:* monitor vlan dog2
config ports 1:* monitor vlan dog3
config ports 1:* monitor vlan dog4
```

After you have configured the ports for monitoring, you can use the show ports vlan statistics command to display information for the configured ports:

```
show ports 1:* vlan statistics
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all platforms.

config protocol add

```
config protocol <protocol_name> add <protocol_type> <hex_value>
{<protocol_type> <hex_value>} ...
```

Description

Configures a user-defined protocol filter.

Syntax Description

protocol_name	Specifies a protocol filter name.
protocol_type	Specifies a protocol type. Supported protocol types include: • `etype` – IEEE Ethertype. • `llc` – LLC Service Advertising Protocol. • `snap` – Ethertype inside an IEEE SNAP packet encapsulation.
hex_value	Specifies a four-digit hexadecimal number between 0 and FFFF that represents: • The Ethernet protocol type taken from a list maintained by the IEEE. • The DSAP/SSAP combination created by concatenating a two-digit LLC Destination SAP (DSAP) and a two-digit LLC Source SAP (SSAP). • The SNAP-encoded Ethernet protocol type.

Default

N / A.

Usage Guidelines

A maximum of 15 protocol filters, each containing a maximum of six protocols, can be defined.

The protocol filter must already exist before you can use this command: use the `create protocol` command to create the protocol filter.

On the *"i"* series platform, all fifteen protocol filters can be active and configured for use. On all other platforms, no more than seven protocols can be active and configured for use.

Example

The following command configures a protocol named Fred by adding protocol type LLC SAP with a value of FFEF:

```
config protocol fred add llc feff
```

History

This command was first available in ExtremeWare 1.0.

Platform Availability

This command is available on all platforms.

config protocol delete

```
config protocol <protocol_name> delete <protocol_type> <hex_value>
{<protocol_type> <hex_value>} ...
```

Description

Deletes the specified protocol type from a protocol filter.

Syntax Description

protocol_name	Specifies a protocol filter name.
protocol_type	Specifies a protocol type. Supported protocol types include: • `etype` – IEEE Ethertype. • `llc` – LLC Service Advertising Protocol. • `snap` – Ethertype inside an IEEE SNAP packet encapsulation.
hex_value	Specifies a four-digit hexadecimal number between 0 and FFFF that represents: • The Ethernet protocol type taken from a list maintained by the IEEE. • The DSAP/SSAP combination created by concatenating a two-digit LLC Destination SAP (DSAP) and a two-digit LLC Source SAP (SSAP). • The SNAP-encoded Ethernet protocol type.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes protocol type LLC SAP with a value of FFEF from protocol *Fred*:

```
config protocol fred delete llc feff
```

History

This command was first available in ExtremeWare 1.0.

Platform Availability

This command is available on all platforms.

config vlan add ports

```
config vlan <name> add ports <portlist> {tagged | untagged} {nobroadcast}
{soft-rate-limit}
```

Description

Adds one or more ports in a VLAN.

Syntax Description

name	Specifies a VLAN name.
portlist	Specifies a list of ports or slots and ports. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
tagged	Specifies the ports should be configured as tagged.
untagged	Specifies the ports should be configured as untagged.
nobroadcast	Prevents broadcasts, multicasts, and unknowns from being transmitted on these ports.
soft-rate-limit	Specifies that these ports should be added as rate-shaped ports. (EW 6.0)

Default

Untagged.

Usage Guidelines

The VLAN must already exists before you can add (or delete) ports: use the `create vlan` command to create the VLAN.

If the VLAN uses 802.1Q tagging, you can specify tagged or untagged port(s). If the VLAN is untagged, the ports cannot be tagged.

Untagged ports can only be a member of a single VLAN. By default, they are members of the default VLAN (named *Default*). In order to add untagged ports to a different VLAN, you must first remove them from the default VLAN. You do not need to do this to add them to another VLAN as tagged ports.

You must configure a loopback port with a unique loopback VLAN tag ID before adding rate-shaped ports.

Example

The following command assigns tagged ports 1, 2, 3, and 6 to a VLAN named *accounting*:

```
config vlan accounting add ports 1, 2, 3, 6 tagged
```

History

This command was first available in ExtremeWare 1.0.

Platform Availability

This command is available on all platforms.

config vlan add ports loopback-vid

```
config vlan <name> add ports <portlist> loopback-vid <vlan-id>
```

Description

Adds a loopback port to a VLAN.

Syntax Description

name	Specifies a VLAN name.
port	Specifies a loopback port for the VLAN.
vlan-id	Specifies a unique loopback VLAN tag.

Default

Untagged.

Usage Guidelines

The VLAN must already exists before you can add (or delete) ports: use the `create vlan` command to create the VLAN.

You must configure a loopback port with a unique loopback VLAN tag ID before adding rate-shaped ports.

Example

The following example sets up bi-directional rate shaping using a loopback port and a rate-shaped port.

First, create the VLAN that will have rate-shaped ports as members:

```
create vlan ratelimit
```

Create the loopback port to rate-shape ingress traffic:

```
config vlan ratelimit add ports 1 loopback-vid 100
```

Configure the user port that will be rate-shaped:

```
config vlan ratelimit add ports 2 soft-rate-limit
```

Configure rate-shaping to be at 5% maximum bandwidth for ingress and egress traffic:

```
config qosprofile QP1 minbw 0 % maxbw 5 % priority low 1,2
```

Enable the loopback port:

```
restart ports 1
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on all *i*-series platforms.

config vlan delete port

```
config vlan <name> delete port <portlist>
```

Description

Deletes one or more ports in a VLAN.

Syntax Description

name	Specifies a VLAN name.
portlist	A list of ports or slots and ports. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

N/A.

Usage Guidelines

None.

Example

The following command removes ports 1, 2, 3, and 6 from a VLAN named *accounting*:

```
config accounting delete port 1, 2, 3, 6
```

History

This command was first available in ExtremeWare 1.0.

Platform Availability

This command is available on all platforms.

config vlan ipaddress

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

Description

Assigns an IP address and an optional subnet mask to the VLAN.

Syntax Description

name	Specifies a VLAN name.
ipaddress	Specifies an IP address.
netmask	Specifies a subnet mask in dotted-quad notation (e.g. 255.255.255.0).
mask length	Specifies a subnet mask as the number of bits (e.g. /24).

Default

N/A.

Usage Guidelines

The VLAN must already exists before you can assign an IP address: use the `create vlan` command to create the VLAN.

NOTE

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

Example

The following commands are equivalent; both assign an IP address of 10.12.123.1 to a VLAN named *accounting*:

```
config vlan accounting ipaddress 10.12.123.1/24
config vlan accounting ipaddress 10.12.123.1 255.255.255.0
```

History

This command was first available in ExtremeWare 1.0.

Platform Availability

This command is available on all platforms.

config vlan name

```
config vlan <old_name> name <new_name>
```

Description

Renames a previously configured VLAN.

Syntax Description

old_name	Specifies the current (old) VLAN name.
new_name	Specifies a new name for the VLAN.

Default

N/A.

Usage Guidelines

You cannot change the name of the default VLAN “Default”

Example

The following command renames VLAN *vlan1* to *engineering*:

```
config vlan vlan1 name engineering
```

History

This command was first available in ExtremeWare 4.1.

Platform Availability

This command is available on all platforms.

config vlan protocol

```
config vlan <name> protocol [<protocol_name> | any]
```

Description

Configures a VLAN to use a specific protocol filter.

Syntax Description

name	Specifies a VLAN name.
protocol_name	Specifies a protocol filter name. This can be the name of a predefined protocol filter, or one you have defined. The following protocol filters are predefined: • IP • IPX • NetBIOS • DECNet • IPX_8022 • IPX_SNAP • AppleTalk any indicates that this VLAN should act as the default VLAN for its member ports.

Default

Protocol Any.

Usage Guidelines

If the keyword any is specified, all packets that cannot be classified into another protocol-based VLAN are assigned to this VLAN as the default for its member ports.

Use the `config protocol` command to define your own protocol filter.

Example

The following command configures a VLAN named accounting as an IP protocol-based VLAN:

```
config accounting protocol ip
```

History

This command was first available in ExtremeWare 1.0.

Platform Availability

This command is available on all platforms.

config vlan tag

```
config vlan <name> tag <vlanid>
```

Description

Assigns a unique 802.1Q tag to the VLAN.

Syntax Description

name	Specifies a VLAN name.
vlanid	Specifies a VLANid to use as an 802.1Q tag. The valid range is from 2 to 4,095.

Default

The default VLAN uses an 802.1Q tag (and an internal VLANid) of 1.

Usage Guidelines

If any of the ports in the VLAN will use an 802.1Q tag, a tag must be assigned to the VLAN. The valid range is from 2 to 4,095 (tag 1 is assigned to the default VLAN).

The 802.1Q tag will also be used as the internal VLANid by the switch.

You can specify a value that is currently used as an internal VLANid on another VLAN; it will become the VLANid for the VLAN you specify, and a new VLANid will be automatically assigned to the other untagged VLAN.

Example

The following command assigns a VLANid of 120 to a VLAN named *accounting*:

```
config accounting tag 120
```

History

This command was first available in ExtremeWare 1.0.

Platform Availability

This command is available on all platforms.

create protocol

```
create protocol <protocol_name>
```

Description

Creates a user-defined protocol filter.

Syntax Description

protocol_name	Specifies a protocol filter name. The protocol filter name can have a maximum of 31 characters.

Usage Guidelines

Protocol-based VLANs enable you to define packet filters that the switch can use as the matching criteria to determine if a particular packet belongs to a particular VLAN.

After you create the protocol, you must configure it using the `config protocol` command. To assign it to a VLAN, use the `config vlan <name> protocol` command.

Example

The following command creates a protocol named *fred*:

```
create protocol fred
```

History

This command was first available in ExtremeWare 1.0.

Platform Availability

This command is available on all platforms.

create vlan

```
create vlan <name>
```

Description

Creates a named VLAN.

Syntax Description

name	Specifies a VLAN name (up to 32 characters).

Default

A VLAN named *Default* exists on all new or initialized Extreme switches:

- It initially contains all ports on a new or initialized switch, except for the management port(s), if there are any.
- It has an 802.1Q tag of 1.
- The default VLAN is untagged on all ports.
- It uses protocol filter any.

An untagged VLAN named *MacVlanDiscover* exists on all new or initialized "*i*" series switches:

- It initially contains no ports.
- It does not initially use an 802.1Q tag, and is assigned the next available internal VLANid starting with 4095.

A VLAN named *Mgmt* exists on switches that have management modules or management ports.

- It initially contains the management port(s) the switch.
- It is assigned the next available internal VLANid as an 802.1Q tag.

Usage Guidelines

A newly-created VLAN has no member ports, is untagged, and uses protocol filter "any" until you configure it otherwise. Use the various `config vlan` commands to configure the VLAN to your needs.

Internal VLANids are assigned automatically using the next available VLANid starting from the high end (4095) of the range.

By default the switch supports 1024 VLANs. The switch can support a maximum of 3000 VLANs if the `CPU-transmit-priority` is set to `normal`, rather than `high` (the default). Use the `config cpu-transmit-priority` command to change the CPU transmit priority (v6.2 or later).

Each VLAN name can be up to 32 standard alphanumeric characters, but 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.

Example

The following command creates a VLAN named *accounting*:

```
create vlan accounting
```

History

This command was first available in ExtremeWare 1.0.

Platform Availability

This command is available on all platforms.

delete protocol

```
delete protocol <protocol_name>
```

Description

Deletes a user-defined protocol.

Syntax Description

protocol_name	Specifies a protocol name.

Default

N/A.

Usage Guidelines

If you delete a protocol that is in use by a VLAN, the protocol associated with than VLAN will become "None."

Example

The following command deletes a protocol named *fred*:

```
delete protocol fred
```

History

This command was first available in ExtremeWare 1.0.

Platform Availability

This command is available on all platforms.

delete vlan

```
delete vlan <name>
```

Description

Deletes a VLAN.

Syntax Description

name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

If you delete a VLAN that has untagged port members, and you want those ports to be returned to the default VLAN, you must add them back explicitly using the `config vlan add port` command.

Example

The following command deletes the VLAN *accounting*:

```
delete accounting
```

History

This command was first available in ExtremeWare 1.0.

Platform Availability

This command is available on all platforms.

disable gvrp

```
disable gvrp
```

Description

Disables the Generic VLAN Registration Protocol (GVRP).

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command globally disables GVRP functionality on the switch. It does not change the GVRP configuration of individual ports, but GVRP will no longer function on these ports.

GVRP is not supported in ExtremeWare version 6.1 or later.

Example

The following command disables GVRP functionality:

```
disable gvrp
```

History

This command was first available in ExtremeWare 2.0.

This command exists but is not supported in ExtremeWare version 6.1 or later.

Platform Availability

This command is available on all platforms.

disable mac-vlan port

```
disable mac-vlan port <portlist>
```

Description

Disables a port from using the MAC-based VLAN algorithm.

Syntax Description

portlist	A list of ports or slots and ports. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

N/A.

Usage Guidelines

Disabling a port removes it from the MacVlanDiscover VLAN. But does not automatically return it to the default VLAN. If you need this port to be a member of the default VLAN, you must explicitly add it back.

Example

The following command disables ports 16 and 17 from using the MAC-based VLAN algorithm:

```
disable mac-vlan port 16,17
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

enable gvrp

```
enable gvrp
```

Description

Enables the Generic VLAN Registration Protocol (GVRP).

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

The GVRP protocol allows switches to automatically discover some of the VLAN information that would otherwise have to be manually configured in each switch.

GVRP must be enabled on individual ports before GVRP information will be sent or received.

By default, GVRP is enabled for both sending and receiving on all ports, so executing this command will normally “turn on” GVRP functionality.

GVRP is not supported in ExtremeWare version 6.1 or later.

Example

The following command enables GVRP functionality:

```
enable gvrp
```

History

This command was first available in ExtremeWare 2.0.

This command exists but is not supported in ExtremeWare version 6.1 or later.

Platform Availability

This command is available on all platforms.

enable mac-vlan mac-group port

```
enable mac-vlan mac-group [any | <group_number>] port <portlist>
```

Description

Enables a port to use the MAC-based VLAN algorithm.

Syntax Description

group_number	A group number that should be associated with a specific set of ports. Specified as an integer.
	any indicates that these ports can be considered members of any MAC group.
portlist	A list of ports or slots and ports. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

N/A.

Usage Guidelines

Enabling ports for MAC-based VLAN usage automatically adds them to the VLAN *MacVlanDiscover* as untagged ports.

In order to enable ports as part of a MAC group, they cannot be untagged members of any other VLAN. Before you can enable them, you must ensure that they have been removed from the default VLAN (named *Default*).

Example

The following set of commands removes ports 16 and 17 form the default VLAN, and then enables them for use with the MAC-based VLAN, associated with any MAC group:

```
config default delete port 16, 17
enable mac-vlan mac-group any port 16,17
```

The following commands enable ports 11 and 12 for use with a MAC-based VLAN, associated with MAC group 10:

```
config default delete port 11, 12
enable mac-vlan mac-group 10 port 11,12
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

show gvrp

```
show gvrp
```

Description

Displays the current configuration and status of GVRP.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

GVRP is not supported in ExtremeWare version 6.1 or later.

Example

The following shows results of this command:

```
GVRP running (866422):   JoinTime 20  LeaveTime 200  LeaveAllTime 1000 cs
GVRP transmit 0  receive 0  tx errors 0  rx errors 0  int errors 0
Enabled for Tx/Rx on ports:     123456789
                10111213141516171819
                20212223242526272829
                303132
VLAN/Ports (t=static tagged, u=static untag, G=GVRP tagged, g=GVRP untag)
 Default (Tag 1)
uuuuuuuuuu..uuu..uuuuuuuuuuuuuuu
 Mgmt (Tag 4094)
.............................
 nat (Tag 4093)
.............................
```

History

This command was first available in ExtremeWare 2.0.

This command exists but is not supported in ExtremeWare version 6.1 or later.

Platform Availability

This command is available on all platforms.

show mac-vlan

```
show mac-vlan {configuration | database}
```

Description

Displays the MAC-based VLAN configuration and MAC address database content.

Syntax Description

configuration	Specifies display of the MAC-based VLAN configuration only.
database	Specifies display of the MAC address database content only.

Default

Shows both configuration and database information.

Usage Guidelines

Use the keyword `configuration` to display only the top section of this information. Use the database keyword to display only the lower section.

Example

The following is an example of the `show mac-vlan` command:

```
Port    Vlan            Group   State
11      MacVlanDiscover 10      Discover
12      MacVlanDiscover 10      Discover
16      MacVlanDiscover any     Discover
17      MacVlanDiscover any     Discover

Total Entries in Database:2
Mac                     Vlan       Group
00:00:00:00:00:AA     anntest1     any
        any           anntest1     10
2 matching entries
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

show protocol

```
show protocol {<protocol>}
```

Description

Displays protocol filter definitions.

Syntax Description

protocol	Specifies a protocol filter name.

Default

Displays all protocol filters.

Usage Guidelines

Displays the defined protocol filter(s) with the types and values of its component protocols.

Example

The following is an example of the `show protocol` command:

```
Protocol Name         Type  Value
----------------      ----- ------
IP                    etype 0x0800
                      etype 0x0806
ipx                   etype 0x8137
netbios                 llc 0xf0f0
                        llc 0xf0f1
decnet                etype 0x6003
                      etype 0x6004
appletalk              snap 0x809b
                       snap 0x80f3
ipx_8022                llc 0xe0e0
ipx_snap               snap 0x8137
```

History

This command was first available in ExtremeWare 1.0.

Platform Availability

This command is available on all platforms.

show vlan

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

Description

Displays information about VLANs.

Syntax Description

name	Specifies a VLAN name.
detail	Specifies that detailed information should be displayed for each VLAN.
stats	Specifies a real-time display of utilization statistics (packets transmitted and received) for a specific VLAN.

Default

Summary information for all VLANs on the device.

Usage Guidelines

Unlike many other vlan-related commands, the keyword “vlan” is required in all forms of this command except when requesting information for a specific vlan.

Use the command `show vlan` to display summary information for all VLANs. It shows various configuration options as a series of “flags” (see the example below). VLAN and protocol names may be abbreviated in this display.

Use the command `show vlan detail` to display detailed information for all VLANs. This displays the same information as for an individual VLAN, but shows every VLAN, one-by-one. After each VLAN display you can elect to continue or quit.

Protocol None indicates that this VLAN was configured with a user-defined protocol that has subsequently been deleted.

Use the command `show vlan stats <name>` to show real-time statistics on the number of packets transmitted and received for the named VLAN. This command will continue to run until you cancel it using the [Esc] key.

Example

The following is an example of the `show vlan` command:

```
Name            VID  Protocol Addr         Flags         Proto        Super Ports
Default         0001 ------------------    -----         ANY          0/ 32
MacVlanDiscover 4095 ------------------    -----         ANY          0/  0
Mgmt            4094 10.205.0.74    /24    -----------   ANY          1/  1
anntest1        4093 ------------------    -----         ANY          0/  0

Flags :  M=ESRP Master, E=ESRP Slave, G=GVRP Enabled, L=Loopback Enabled
        S=SuperVlan, s=SubVlan, R=SubVLAN IP Range Configured
        C=Domain-masterVlan, c=Domain-memberVlan
        f=IP Forwarding Enabled, m=IPmc Forwarding Enabled
        r=RIP Enabled, o=OSPF Enabled, p=PIM Enabled, d=DVMRP Enabled
        R=IPX RIP Enabled, P=IPX SAP Enabled
        N=GNS Reply Enabled, 2=IPX Type 20 Forwarding Enabled
```

The following is an example of the `show vlan Default` command:

```
VLAN Interface[0-200] with name "Default" created by user
     Tagging:  802.1Q Tag 1
     STPD:     Domain "s0" is not running spanning tree protocol
     Protocol: Match all unfiltered protocols.
     Loopback: Disable
     Rate Shape:   Disable
     QosProfile:   QP1
     Ports:   32.     (Number of active ports=0)
               Untag:  1  2  3  4  5  6  7  8  9  10  11  12  13  14
               15  16  17  18  19  20  21  22  23  24  25  26  27  28  29
              30  31  32
```

History

This command was first available in ExtremeWare 1.0.

This command was modified to support longer VLAN names in ExtremeWare 6.2.2.

Platform Availability

This command is available on all platforms.

unconfig ports monitor vlan

```
unconfig ports <portlist> monitor vlan <name>
```

Description

Removes port-based VLAN monitoring.

Syntax Description

portlist	Specifies one or more ports. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

None.

Example

The following command removes monitoring for ports on VLAN *accounting*:

```
unconfig ports 8:1-8:6 monitor vlan accounting
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all platforms.

unconfig vlan ipaddress

```
unconfig vlan <name> ipaddress
```

Description

Removes the IP address of the VLAN.

Syntax Description

name	Specifies a VLAN name.
ipaddress	Specifies that the ipaddress association with this VLAN should be cleared.

Default

N/A.

Usage Guidelines

None.

Example

The following command removes the IP address from the VLAN *accounting*:

```
unconfig vlan accounting ipaddress
```

History

This command was first available in ExtremeWare 1.0.

Platform Availability

This command is available on all platforms.

6 FDB Commands

This chapter describes commands for:

- Configuring FDB entries
- Displaying FDB entries

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.

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.

The FDB has four types of entries:

- 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 of 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.
- 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 create permanent entries. A permanent entry can either be a unicast or multicast MAC address. All entries entered through the command line interface (CLI) are stored as permanent. The stand-alone switches can support a maximum of 64 permanent entries, and the modular switches support a maximum of 254 permanent entries.
- Blackhole entries — A blackhole entry configures the switch to discard packets with a specified MAC destination address. 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.

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 CLI.

A QoS profile can be associated 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.

clear fdb

```
clear fdb {<mac_address> | broadcast-mac | locked-mac <mac_address> | vlan
<name> | ports <portlist>}
```

Description

Clears dynamic FDB entries that match the filter.

Syntax Description

mac_address	Specifies a MAC address, using colon-separated bytes.
broadcast-mac	Specifies the broadcast MAC address. May be used as an alternate to the colon-separated byte form of the address ff:ff:ff:ff:ff:ff.
locked-mac <mac_address>	Specifies the MAC address of a locked static FDB entry.
name	Specifies a VLAN name.
portlist	Specifies one or more ports. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

Clears all FDB entries.

Usage Guidelines

This command clears FDB entries based on the specified criteria. When no options are specified, the command clears all FDB entries.

Example

The following two commands are synonymous, and clear FDB entries for the broadcast MAC address:

```
clear fdb broadcast-mac
clear fdb ff:ff:ff:ff:ff:ff
```

The following command clears any locked-static FDB entries for MAC address 00:E0:2B:00:00:00:

```
clear fdb locked-mac 00:E0:2B:00:00:00
```

The following command clears any FDB entries associated with ports 3-5:

```
clear fdb ports 3-5
```

The following command clears any FDB entries associated with VLAN *corporate*:

```
clear fdb vlan corporate
```

History

This command was available in ExtremeWare 2.0.

The command was modified in ExtremeWare 6.2.1 to support the broadcast-mac keyword and to support clearing locked-static entries.

Platform Availability

This command is available on all platforms.

config fdb agingtime

```
config fdb agingtime <seconds>
```

Description

Configures the FDB aging time for dynamic entries.

Syntax Description

seconds	Specifies the aging time in seconds. Range is 15 through 1,000,000. A value of 0 indicates that the entry should never be aged out.

Default

300 seconds.

Usage Guidelines

The range is 15 through 1,000,000 seconds.

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 out, but non-permanent static entries can be deleted if the switch is reset.

Example

The following command sets the FDB aging time to 3,000 seconds:

```
config fdb agingtime 3000
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

create fdbentry vlan blackhole

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

Description

Creates a blackhole FDB entry.

Syntax Description

mac_address	Specifies a device MAC address, using colon-separated bytes.
name	Specifies a VLAN name associated with a MAC address.
blackhole	Configures the MAC address as a blackhole entry.
source-mac	Specifies that the blackhole MAC address matches the ingress source MAC address. Support for this parameter was added in ExtremeWare 6.2.
dest-mac	Specifies that the blackhole MAC address matches the egress destination MAC address. Support for this parameter was added in ExtremeWare 6.2.
both	Specifies that the blackhole MAC address matches the ingress source MAC address or the egress destination MAC address. Support for this parameter was added in ExtremeWare 6.2.

Default

N/A.

Usage Guidelines

Blackhole entries are useful as a security measure or in special circumstances where packets with a specific source or destination address must be discarded.

A blackhole entry configures the switch to discard packets with the specified MAC address. You can specify whether the MAC address should match the source (ingress) MAC address, or the destination (egress) MAC address, or both.

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. In the output from a `show fdb` command, entries will have “p” flag (permanent) set, as well as the “b” (for ingress blackhole) and/or “B” (for egress blackhole) flags set.

Example

The following example adds a blackhole entry to the FDB for MAC address is 00 E0 2B 12 34 56, in VLAN *marketing* on port 4:

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

History

This command was available in ExtremeWare 2.0.

Support for specifying source or destination MAC address was added in ExtremeWare 6.2.

Platform Availability

This command is available on all platforms.

create fdbentry vlan dynamic

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

Description

Creates a permanent dynamic FDB entry, and associates it with an ingress and/or egress QoS profile.

Syntax Description

mac_address	Specifies a device MAC address, using colon separated bytes.
broadcast-mac	Specifies the broadcast MAC address. May be used as an alternate to the colon-separated byte form of the address ff:ff:ff:ff:ff:ff.
any-mac	Specifies the wildcard, permanent FDB entry used to give higher priority to an 802.1p packet.
vlan <name>	Specifies a VLAN name associated with a MAC address.
dynamic	Specifies that the entry will be learned dynamically.
qosprofile	QoS profile associated with the destination MAC address of the egress port.
inqosprofile	QoS profile associated with the source MAC address of the ingress port. Support for this parameter was added in ExtremeWare 6.2.

Default

N/A.

Usage Guidelines

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

The FDB entry is not actually created until the MAC address is encountered as the source MAC address in a packet. Thus, initially the entry may not appear in the `show fdb` output. Once the entry has been learned, it is created as a permanent dynamic entry, designated by “dpm” in the flags field of the `show fdb` output.

A dynamic entry is flushed and relearned (updated) when any of the following take place:

- A VLAN is deleted.
- A VLAN identifier (VLANid) is changed.
- A port mode is changed (tagged/untagged).
- A port is deleted from a VLAN.
- A port is disabled.
- A port enters blocking state.
- A port QoS setting is changed.

- A port goes down (link down).

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

You can use the `show fdb permanent` command to display permanent FDB entries, including their QoS profile associations.

Example

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

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

QoS profile *qp2* will be applied when the entry is learned.

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

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

History

This command was available in ExtremeWare 2.0.

Support for associating separate QoS profiles with ingress and egress ports was added in ExtremeWare 6.2.

This command was modified in ExtremeWare 6.2.1 to support the `broadcast-mac` option.

Platform Availability

This command is available on all platforms.

create fdbentry vlan ports

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

Description

Creates a permanent static FDB entry, and optionally associates it with an ingress and/or egress QoS profile.

Syntax Description

mac_address	Specifies a device MAC address, using colon-separated bytes.
name	Specifies a VLAN name associated with a MAC address.
portlist	Specifies one or more ports associated with the MAC address. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
qosprofile	QoS profile associated with the destination MAC address of the egress port
inqosprofile	QoS profile associated with the source MAC address of the ingress port. Support for this parameter was added in ExtremeWare 6.2.

Default

N/A.

Usage Guidelines

If more than one port number is associated with a permanent MAC entry, packets are multicast to the multiple destinations.

Permanent entries are retained in the database if the switch is reset or a power off/on cycle occurs. A permanent static entry can either be a unicast or multicast MAC address. The stand-alone switches can support a maximum of 64 permanent entries, and the modular switches support a maximum of 254 permanent entries.

After they have been created, permanent static entries stay the same as when they were created. If the same MAC address is encountered on another virtual port that is not included in the permanent MAC entry, it is handled as a blackhole entry. The static entry 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).

Permanent static entries are designated by "spm" in the flags field of the `show fdb` output. You can use the `show fdb permanent` command to display permanent FDB entries, including their QoS profile associations.

Example

The following example adds a permanent, static entry to the FDB for MAC address is 00 E0 2B 12 34 56, in VLAN *marketing* on port 4:

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

History

This command was available in ExtremeWare 2.0.

Support for associating separate QoS profiles with ingress and egress ports was added in ExtremeWare 6.2.

Platform Availability

This command is available on all platforms.

delete fdbentry

```
delete fdbentry [[<mac_address> | broadcast-mac] vlan <name> | all]
```

Description

Deletes one or all permanent FDB entries.

Syntax Description

mac_address	Specifies a device MAC address, using colon-separated bytes.
broadcast-mac	Specifies the broadcast MAC address. May be used as an alternate to the colon-separated byte form of the address ff:ff:ff:ff:ff:ff.
name	Specifies a VLAN name.
all	Specifies that all FDB entries should be deleted.

Default

N/A.

Usage Guidelines

None.

Example

The following example deletes a permanent entry from the FDB:

```
delete fdbentry 00:E0:2B:12:34:56 vlan marketing
```

The following example deletes all permanent entry from the FDB:

```
delete fdbentry all
```

History

This command was available in ExtremeWare 2.0.

This command was modified in ExtremeWare 6.2.0 to support the `all` option.

This command was modified in ExtremeWare 6.2.1 to support the `broadcast-mac` option.

Platform Availability

This command is available on all platforms.

run fdb-check

```
run fdb-check [index <bucket> <entry> | [<mac_address> | broadcast-mac]
{<vlan-name>}] {extended} {detail}
```

Description

Checks MAC FDB entries for consistency.

Syntax Description

bucket	Specifies the bucket portion of the FDB hash index.
entry	Specifies the entry portion of the FDB hash index.
mac-address	Specifies a MAC address (hex octet). FDB entries with this MAC address will be checked.
broadcast-mac	Specifies the broadcast MAC address. May be used as an alternate to the hex octet form, ff:ff:ff:ff:ff:ff. (6.2.1 and higher)
vlan-name	Specifies a VLAN name. FDB entries for this VLAN with the specified MAC address will be checked.
extended	Enables OTP index checking in the MAC entry and VPST of the egress port.
detail	Specifies that more detailed debug information should be logged.

Default

N / A.

Usage Guidelines

The FDB error checking function logs the error count into the system log. Use the `detail` option to log more detailed debug information.

If you do not enter a VLAN name, ExtremeWare check all FDB entries with the specified MAC address.

Example

Given the following FDB entry on an MSM 64:

```
Index              Mac              Vlan       Age  Use  Flags Port List

-------------------------------------------------------------------------------
  cf3c0-006 00:00:00:00:00:01         v1(4093)  0540 0000 d m      3:4
```

All the following commands will do consistency checking on this entry:

```
run fdb-check 00:00:00:00:00:01
run fdb-check 00:00:00:00:00:01 detail
run fdb-check 00:00:00:00:00:01 extended detail
run fdb-check 00:00:00:00:00:01 vlan v1
run fdb-check index cf3c 0 extended detail
```

History

This command was first available in ExtremeWare 6.1.9

This command was modified in ExtremeWare 6.2.1 to support the `broadcast-mac` keyword.

Platform Availability

This command is available on all *"i"* series platforms.

The `extended` option is available on the Black Diamond 6800 chassis-based system only.

show fdb

```
show fdb {<mac_address> | broadcast-mac | vlan <name> | <portlist> |
permanent}
```

Description

Displays FDB entries.

Syntax Description

mac_address	Specifies a MAC address, using colon-separated bytes, for which FDB entries should be displayed.
broadcast-mac	Specifies the broadcast MAC address. May be used as an alternate to the colon-separated byte form of the address ff:ff:ff:ff:ff:ff.
name	Displays the entries for a specific VLAN.
portlist	Displays the entries for one or more ports. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
permanent	Displays all permanent entries, including the ingress and egress QoS profiles.

Default

All.

Usage Guidelines

Displays FDB entries as specified, or displays all FDB entries.

The show output displays the following information:

EQP	The Ingress QoS profile assigned to the entry (appears only if the keyword permanent is specified)
IQP	The Egress QoS profile assigned to the entry (appears only if the keyword permanent is specified)
Index	The FDB hash index, in the format <bucket>-<entry>
Mac	The MAC address that defines the entry.
Vlan	The VLAN for the entry
Age	The age of the entry, in seconds
Use	

Flags	Flags that define the type of entry: • d - Dynamic • s - Static • p - Permanent • m - MAC • i - an entry also exists in the IP FDB • x - an entry also exists in the IPX FDB • l - lockdown MAC entry • M - Mirror • B - Egress Blackhole • b - Ingress Blackhole
Port List	The ports on which the MAC address has been learned

Example

The following command displays information about all the entries in the FDB:

```
show fdb
```

It produces output similar to the following:

```
Index              Mac              Vlan      Age  Use  Flags  Port List
-------------------------------------------------------------------------
0a0e0-100 00:01:30:EC:D3:00        lab(4000)  0000 0001 d i      1
2b560-ffb 01:00:0C:CC:CC:CD           (0000)  0000 0000 s m      CPU
30040-ffb 00:E0:2B:00:00:00        zzz(0652)  0000 0000 s m      CPU
332890-ffb 00:E0:2B:00:00:00   Default(0001)  0000 0000 s m      CPU
3d760-ffb 00:E0:2B:00:00:00       Mgmt(4094)  0000 0000 s m      CPU
3d770-ffb 00:E0:2B:00:00:00 MacVlanDis(4095)  0000 0000 s m      CPU
42560-ff0 00:01:30:6C:0D:00        lab(4000)  0000 0000 s m      CPU
46460-100 00:10:E3:1D:00:1E        lab(4000)  0000 0001 d i      1
4d060-100 00:10:E3:1D:00:05        lab(4000)  0000 0001 d i      1
4df70-ff0 00:01:30:6C:0D:00    Default(0001)  0000 0000 s m      CPU
4f7a0-ff0 00:01:30:6C:0D:00        zzz(0652)  0000 0000 s m      CPU
51f50-100 00:01:30:CA:F6:00        lab(4000)  0000 0001 d i      1
        . . .
67b20-100 00:30:D3:01:5A:E0        lab(4000)  0000 0001 d i      1
80a10-204 FF:FF:FF:FF:FF:FF        lab(4000)  0000 0000 s m      CPU, 2, 1
80fe0-208 FF:FF:FF:FF:FF:FF MacVlanDis(4095)  0000 0000 s m      CPU
80ff0-202 FF:FF:FF:FF:FF:FF       Mgmt(4094)  0000 0000 s m      CPU
8d8d0-20a FF:FF:FF:FF:FF:FF        zzz(0652)  0000 0000 s m      CPU, 2
8f000-200 FF:FF:FF:FF:FF:FF    Default(0001)  0000 0000 s m      CPU
98670-100 00:01:30:E7:F2:00        lab(4000)  0000 0001 d i      1
fcf70-202 00:E0:2B:00:00:02       Mgmt(4094)  0000 0000 s m      CPU

Flags : d - Dynamic, s - Static, p - Permanent, m - MAC, i - IP,
        x - IPX, l - lockdown MAC, M - Mirror, B - Egress Blackhole,
        b - Ingress Blackhole.

Total: 33 Static: 16 Perm: 0 Locked: 0 Dyn: 17 Dropped: 0
FDB Aging time: 300 seconds
```

The following command displays information about the permanent entries in the FDB:

```
show fdb permanent
```

It produces output similar to the following:

```
EQP IQP Index              Mac              Vlan       Age  Use  Flags Port List
-------------------------------------------------------------------------------
QP3 QP2 ----- --- 00:10:E3:1D:00:05   anntest1(4094)  ---- ----  pm     ---
QP3 QP2 4e610-206 00:01:03:2F:38:EE   anntest1(4094)  0000 0000 spm       ---
QP3 QP2 ----- --- 00:60:B0:F9:58:9D    Default(0001)  ---- ----  pm     ---

Flags : d - Dynamic, s - Static, p - Permanent, m - MAC, i - IP,
         x - IPX, l - lockdown MAC, M - Mirror, B - Egress Blackhole,
         b - Ingress Blackhole.
```

History

This command was available in ExtremeWare 2.0.

This command was modified in ExtremeWare 6.2.1 to support the `broadcast-mac` option.

Platform Availability

This command is available on all platforms.

7 QoS Commands

This chapter describes the following commands:

- Commands for configuring Quality of Service (QoS) profiles
- Commands creating traffic groupings and assigning the groups to QoS profiles
- Commands for configuring, enabling and disabling explicit class-of-service traffic groupings (802.1p and Diffserv)
- Commands for configuring, enabling and disabling Random Early Detection (RED)
- Commands for configuring traffic grouping priorities
- Commands for verifying configuration and performance
- Commands for enabling and disabling the Dynamic Link Context System (DLCS)

Qualify of Service (QoS) is a feature of ExtremeWare that allows you to specify different service levels for outbound and inbound traffic. QoS is an effective control mechanism for networks that have heterogeneous traffic patterns. Using QoS, you can specify the service that a traffic type receives.

Policy-based QoS allows you to protect bandwidth for important categories of applications or specifically limit the bandwidth associated with less critical traffic. The switch contains separate hardware queues on every physical port. Each hardware queue is programmed by ExtremeWare with bandwidth management and prioritization parameters, defined as a QoS profile. The bandwidth management 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. Switch products with the "*i*" chipset can be configured with up to eight physical queues per port, while other Extreme switches can be configured with up to four physical queues per port.

Policy-based QoS can be configured to perform per-port Random Early Detection (RED). Using this capability, the switch detects when traffic is filling up in any of the eight hardware queues, and performs a random discard on subsequent packets, based on the configured RED drop-probability. Instead of dropping sessions during times when the queue depth is exceeded, RED causes the switch to lower session throughput. Only switches and modules with the "i" chipset can use RED.

To configure QoS, you define how your switch responds to different categories of traffic by creating and configuring QoS profiles. The service that a particular type of traffic receives is determined by assigning a QoS profile to a traffic grouping or classification. The building blocks are defined as follows:

- QoS profile — Defines bandwidth and prioritization parameters.
- Traffic grouping — A method of classifying or grouping traffic that has one or more attributes in common.

- QoS policy — The combination that results from assigning a QoS profile to a traffic grouping.

QoS profiles are assigned to traffic groupings to modify switch-forwarding behavior. When assigned to a traffic grouping, the combination of the traffic grouping and the QoS profile comprise an example of a single policy that is part of Policy-Based QoS.

Extreme switch products support explicit Class of Service traffic groupings. This category of traffic groupings describes what is sometimes referred to as *explicit packet marking*, and includes:

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

Extreme products that use the "*i*" chipset support DiffServ capabilities. All Extreme switches support the standard 802.1p priority bits that are part of a tagged Ethernet packet.

DLCS

The Dynamic Link Context System (DLCS) is a feature of ExtremeWare and Extreme switches that snoops Windows Internet Naming Service (WINS) NetBIOS packets and creates a mapping between a user name, the IP address or MAC address of the workstation, and a port on the switch. Based on the information in the packet, DLCS can detect when an workstation boots up or a user logs in or out, and dynamically maps the user or workstation name to the current IP address and switch port. For DLCS to operate within ExtremeWare, the user or workstation must allow for automatic DLCS updates.

Information obtained through DLCS is used by the Grouping Manager module found in the EPICenter 3.1 software, and enables the configuration of policies that apply to named users or workstations. Enabling the DLCS feature is only useful if you plan to use the EPICenter software. Currently, there are no other features that can make use of the information that the DLCS feature provides.

clear dlcs

```
clear dlcs
```

Description

Clears all learned DLCS data.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

If the IP address of an end-station changes, and the end-station is not immediately rebooted, the old host-to-IP mapping is not deleted. You must delete the mapping through the ExtremeWare Enterprise Manager Policy System.

Example

The following command clears all learned DLCS data from the switch:

```
clear dlcs
```

History

This command was available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

config diffserv examination code-point qosprofile ports

```
config diffserv examination code-point <code_point> qosprofile <qosprofile>
ports [<portlist> | all] {low-drop-probability | high-drop-probability}
```

Description

Configures the default ingress Diffserv code points (DSCP) to QoS profile mapping.

Syntax Description

code_point	Specifies a DiffServ code point (a 6-bit value in the IP-TOS byte in the IP header).
qosprofile	Specifies the QoS profile to which the Diffserv code point is mapped.
portlist	Specifies a list of ports or slots and ports. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
all	Specifies that this applies to all ports on the device.
low-drop-probability	Specifies that the DSCP has a low drop-probability level. Supported only for SONET ports on a PoS module.
high-drop-probability	Specifies that the DSCP has a high drop-probability level. Supported only for SONET ports on a PoS module.

Default

See Table 8.

Usage Guidelines

You can specify up to 64 different code points for each port. Code point values are grouped and assigned to the default QoS profiles as follows:

Table 8: 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

The mapping is applied in the ingress direction.

The `low-drop-probability` and `high-drop-probability` keywords are applicable only to SONET ports. The `low-drop-probability` and `high-drop-probability` keywords are useful in conjunction with the weighted RED (WRED) implementation provided by SONET ports. This implementation

supports two different drop probabilities; one for DSCPs designated as having low drop-probability and another for DSCPs designated as having high drop-probability. These keywords enable complete flexibility in assigning DSCPs to the two different drop-probability levels.

Example

The following command specifies that packets arriving on ports 5-8 that use code point 25 be assigned to qp2:

```
config diffserv examination code-point 25 qosprofile qp2 ports 5-8
```

The following command sets up the mapping for the EF PHB (PoS module only):

```
config diffserv examination code-point 46 qosprofile qp8 ports 2:1-2:2
```

History

This command was first available in ExtremeWare 6.0.

This command was modified in an ExtremeWare IP Technology Services Release based on v6.1.5b20 to support PoS modules.

Platform Availability

This command is available on the "*i*" series platforms. The PoS module extensions are supported on the BlackDiamond 6800 series chassis-based system only.

config diffserv replacement priority

```
config diffserv replacement priority <value> code-point <code_point> ports
[<portlist> | all]
```

Description

Configures the default egress Diffserv replacement mapping.

Syntax Description

value	Specifies the 802.1p priority value.
code_point	Specifies a 6-bit value to be used as the replacement code point in the IP-TOS byte in the IP header.
portlist	Specifies a list of egress ports or slots and ports. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

N / A.

Usage Guidelines

To replace DiffServ code points you must enable both 802.1p replacement and DiffServ replacement using the `enable dot1p replacement ports` and `enable diffserv replacement ports` commands.

The default 802.1p priority value to code point mappings are described as follows:

Table 9: Default 802.1p Priority Value-to-Code Point Mapping

Hardware Queue "*i*" Chipset	802.1p Priority value	Code Point
Q0	0	0
Q1	1	8
Q2	2	16
Q3	3	24
Q4	4	32
Q5	5	40
Q6	6	48
Q7	7	56

Example

The following command specifies that a code point value of 25 should be used to replace the TOS bits in packets with an 802.1p priority of 2 for ports 5-9:

```
config diffserv replacement priority 2 code-point 25 ports 5-9
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

config dot1p type

```
config dot1p type <dot1p_priority> qosprofile <qosprofile>
```

Description

Configures the default QoS profile to 802.1p priority mapping.

Syntax Description

dot1p_priority	Specifies the 802.1p priority value. The value is an integer between 0 and 7.
qosprofile	Specifies a QoS profile.

Default

N / A.

Usage Guidelines

An 802.1p priority value seen on ingress can be mapped to a particular QoS profile and with specific bandwidth management and priority behavior.

The default mapping of each 802.1p priority value to QoS profile is as follows:

Table 10: 802.1p Priority Value-to-QoS Profile Default Mapping

Priority Value	QoS Profile Summit Chipset	QoS Profile *"i"* Chipset
0	Qp1	Qp1
1	Qp1	Qp2
2	Qp2	Qp3
3	Qp2	Qp4
4	Qp3	Qp5
5	Qp3	Qp6
6	Qp4	Qp7
7	Qp4	Qp8

Example

The following commands swap the QoS profiles associated with 802.1p priority values 1 and 2 on an "*i*" series device:

```
config dot1p type 2 qosprofile qp2
config dot1p type 1 qosprofile qp3
```

History

This command was available in ExtremeWare 6.0.

Platform Availability

This command is available on all platforms.

config ipqos add

```
config ipqos add {tcp | udp | other | all} <dest_ipaddress>/<mask_length>
{l4-dstport <tcp/udp_port_number>} {<ip_src_address>/<mask_length>}
{l4-srcport <tcp/udp_port_number>} [qosprofile <qosname> | blackhole]
```

Description

Applies a set of destination IP addresses to an IPQoS traffic grouping by specifying a network address and subnet mask.

Syntax Description

tcp	Specifies that the TCP protocol is to be used for this traffic grouping.
udp	Specifies that the UDP protocol is used for this traffic grouping.
other	Specifies that an IP protocol other than TCP or UDP is to be used for this traffic grouping.
all	Specifies that any IP protocol may be used for this traffic grouping.
dest_ipaddress/mask_length	The destination IP address (or group of IP addresses) to which the QoS profile is applied. Mask length is a value that indicates the length of the subnet mask to be applied.
ltcp/upd_port_number	The layer 4 destination port number. This is the IP port number associated with the protocol specified in the command string. • If TCP is used as the protocol, the layer 4 port number is a TCP port number. • If UDP is used, the layer 4 port number is a UDP port number. • If not specified, all port numbers used by the IP protocol (TCP or UDP) are implied.
ip_src_address/mask_length	The source IP address (or group of IP addresses) to which the QoS profile is applied. Mask length is a value that indicates the length of the subnet mask to be applied.
ltcp/udp_port_number	The layer 4 source port number. This is the IP port number associated with the protocol specified in the command string. • If TCP is used as the protocol, the layer 4 port number is a TCP port. • If UDP is used, the layer 4 port number is a UDP port.
qosname	Specifies the name of the QoS profile to be used by this traffic grouping.
blackhole	Specifies the blackhole profile.

Default

N/A.

Usage Guidelines

The longer form (also known as a *flow*) is used for specifying additional components of IP packets, such as IP source address, and destination or source TCP/UDP port information.

A long form multicast and unicast entry (flow) has higher precedence over a matching short form multicast and unicast entry (non-flow).

Within the IPQoS long form (flow), precedence is determined by the traffic grouping information provided. For example, an IP QoS policy that includes a specified source IP address has higher precedence than an IP QoS policy that includes a layer 4 source port (but no source IP address). An IP QoS policy containing a layer 4 destination port (but no source IP or layer 4 port number) has the lowest precedence.

When using the `config ipqos` command, the following rules apply:

- The short form of the command only accepts a unicast `<dest_ipaddr>`.
- An IP addr of 0.0.0.0 /0 can be used as a wildcard unicast destination.
- Unless the IntraSubnet QoS (ISQ) feature is enabled, the traffic groupings defined within IPQoS apply to traffic being routed (not layer 2 switched) to the destination IPQoS traffic grouping within the switch.
- IPQoS does not apply to traffic that is normally handled by the switch management processor, including ICMP traffic and packets associated with routing protocols such as OSPF, RIP, DVMRP, and so on.
- Traffic groupings on source IP addresses may utilize a variable subnet mask when an IP multicast destination is specified, but must be a wildcard or specific destination (32 bits of mask) if an IP unicast destination is specified.
- If you are defining a grouping within IPQoS, and you are using the `other` protocol option, the switch filters on the 32 bits after the IP header.
- If you are defining a grouping within IPQoS, and you are using the `all` protocol option, the switch creates three groupings: one grouping for TCP, one grouping for UDP, and one grouping for `other`.
- The IPQoS policies are programmed when a station is added to the forwarding database (FDB). If the station already exists in the IP forwarding database (IPFDB), clear it so that it may be added again using the `clear ipfdb all` command.

Example

The following short-form command defines a traffic grouping for traffic destined to the 10.1.2.X network and assigns it to the *qp2* QoS profile:

```
config ipqos add 10.1.2.3/24 qosprofile qp2
```

The following long-form command groups all traffic to TCP destination port 80 destined for the 10.1.2.x network from 10.1.1.1 using TCP source port 20 and assigns it to *qp4*:

```
config ipqos add tcp 10.1.2.3/24 l4-dstport 80 10.1.1.1/32 l4-srcport 20 qosprofile
qp4
```

History

The short form of this command was available in ExtremeWare 2.0. The long form was available in ExtremeWare 4.0.

Support for this command was discontinued in ExtremeWare 6.0.

Platform Availability

This command is available only on platforms based on the Summit chipset.

config ipqos delete

```
config ipqos delete {tcp | udp | other | all}
<dest_ipaddress>/<mask_length> {l4-dstport <tcp/udp_port_number>}
{<ip_src_address>/<mask_length>} {l4-srcport <tcp/udp_port_number>}
[qosprofile <qosname> | blackhole]
```

Description

Removes a set of destination IP addresses from an IPQoS traffic grouping.

Syntax Description

tcp	Specifies that the TCP protocol is to be used for this traffic grouping.
udp	Specifies that the UDP protocol is used for this traffic grouping.
other	Specifies that an IP protocol other than TCP or UDP is to be used for this traffic grouping.
all	Specifies that any IP protocol may be used for this traffic grouping.
dest_ipaddress/mask_length	The destination IP address (or group of IP addresses) to which the QoS profile is applied. Mask length is a value that indicates the length of the subnet mask to be applied.
tcp/upd_port_number	The layer 4 destination port number. This is the IP port number associated with the protocol specified in the command string. • If TCP is used as the protocol, the layer 4 port number is a TCP port number. • If UDP is used, the layer 4 port number is a UDP port number. • If not specified, all port numbers used by the IP protocol (TCP or UDP) are implied.
ip_src_address/mask_length	The source IP address (or group of IP addresses) to which the QoS profile is applied. Mask length is a value that indicates the length of the subnet mask to be applied.
tcp/udp_port_number	The layer 4 source port number. This is the IP port number associated with the protocol specified in the command string. • If TCP is used as the protocol, the layer 4 port number is a TCP port. • If UDP is used, the layer 4 port number is a UDP port.
qosname	Specifies the name of the QoS profile to be used by this traffic grouping.
blackhole	Specifies the blackhole profile.

Default

N/A.

Usage Guidelines

None.

Example

The following short-form command removes a traffic grouping definition for traffic destined to the 10.1.2.X network:

```
config ipqos delete 10.1.2.3/24 qosprofile qp2
```

The following long-form command removes the traffic group defined as all traffic to TCP destination port 80 destined for the 10.1.2.x network from 10.1.1.1 using TCP source port 20:

```
config ipqos delete tcp 10.1.2.3/24 l4-dstport 80 10.1.1.1/32 l4-srcport 20 qosprofile
qp4
```

History

The short form of this command was available in ExtremeWare 2.0. The long form was available in ExtremeWare 4.0.

Support for this command was superseded in ExtremeWare 6.0 by the `create access-list` command.

Platform Availability

This command is available only on platforms based on the Summit chipset.

config ports qosprofile

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

Description

Configures one or more ports to use a particular QoS profile.

Syntax Description

portlist	Specifies a list of ports or slots and ports. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
qosprofile	Specifies a QoS profile.

Default

N / A.

Usage Guidelines

"i" series switches support eight QoS profiles (QP1 - QP8). Switches based on the Summit chipset support four profiles (QP1 - QP4).

Example

The following command configures port five to use QoS profile QP3:

```
config ports 5 qosprofile QP3
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the *"i"* series platforms.

config qosmode

```
config qosmode [ingress | egress]
```

Description

Changes the QoS mode to ingress mode or egress mode.

Syntax Description

ingress	Specifies ingress mode.
egress	Specifies egress mode.

Default

Ingress.

Usage Guidelines

If you change the QoS mode setting from the default, you must save and reboot the switch in order for the changes to take effect.

You can verify the QoS mode settings by using the `show switch` command.

Example

The following command changes the QoS mode setting from the default setting of ingress:

```
config qosmode egress
```

History

This command was available in ExtremeWare 2.0.

Support for this command was superseded in ExtremeWare 6.0 by the `create access-list` command.

Platform Availability

This command is available only on platforms based on the Summit chipset.

config qosprofile

```
config qosprofile <qosprofile> minbw <min_percent> maxbw <max_percent>
priority <level> {[minbuf <percent> maxbuf <number> [K | M] | maxbuff
<number> [K | M] | <portlist>]}
```

Description

Modifies the default QoS profile parameters.

Syntax Description

qosprofile	Specifies a QoS profile name.
min_percent	Specifies a minimum bandwidth percentage for this queue. The default setting is 0.
max_percent	Specifies the maximum bandwidth percentage this queue is permitted to use. The default setting is 100.
level	Specifies a service priority setting. Settings include low, lowHi, normal, normalHi, medium, mediumHi, high, and highHi. The default setting is low. Available in egress mode only.
percent	Specifies the minimum percentage of the buffer set aside for the queue. Cumulative % of the queues should not exceed 100%.
number	Specifies the maximum buffer size in either M or K bytes. The range is 0 to 16384. The default is 256 K. You must reboot for changes to take effect. • K indicates the value is in K bytes. • M indicates the value is in M bytes.
portlist	Specifies a list of ports or slots and ports to which the parameters apply. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

- Minimum bandwidth — 0%
- Maximum bandwidth — 100%
- Priority — low
- Minimum buffer percent — 0%
- Maximum buffer size — 256K

Usage Guidelines

On Summit chipset-based switches in ingress mode, any changes to parameters of the four predefined QoS profiles have the corresponding effect on the ports to which they are mapped.

The `minbuf` parameter reserves buffer memory for use exclusively by a QoS profile across all affected ports. The default value for buffer settings is 0%. The sum of all QoS profile buffer parameters should not exceed 100%. The `maxbuf` parameter allows you to set a maximum buffer for each queue, so that a single queue will not consume all of the unallocated buffer space. You should not modify the buffer parameter unless specific situations and application behavior indicate. You must reboot the switch for changes to this parameter to take effect.

For ExtremeWare 4.0:

- Only four priority levels are available (low, normal, medium, and high).

Example

The following command configures the parameters of QoS profile *qp1* on a Summit chipset-based switch:

```
config qosprofile qp1 minbw 10% maxbw 80% priority high
```

The following command configures the QoS profile parameters of QoS profile *qp5* for specific ports on "*i*" series switch:

```
config qosprofile qp5 minbw 10% maxbw 80% priority highHi ports 5-7
```

The following command configures the buffer size for QoS profile *qp5* on an "*i*" series switch:

```
config qosprofile qp5 minbw 10% maxbw 80% priority highhi minbuf 3% maxbuff 1024K
```

History

This command was available in ExtremeWare 2.0.

The minbuff, maxbuff, and ports arguments were available in ExtremeWare 6.

Platform Availability

The basic command is available on all platforms.

The minbuff, maxbuff, and ports arguments are available on "*i*" series platforms.

config qostype priority

```
config qostype priority [source-mac | dest-mac | access-list | vlan |
diffserv | dot1p] <priority>
```

Description

Configures the priority of the specified QoS traffic grouping.

Syntax Description

source-mac	Specifies the priority of traffic groupings based on FDB source MAC addresses. Default is 7.
dest-mac	Specifies the priority of traffic groupings based on FDB destination MAC addresses. Default is 8.
access-list	Specifies the priority of access-list based traffic groupings. Default is 11.
vlan	Specifies the priority of VLAN-based traffic groupings. Default is 1.
diffserv	Specifies the priority of traffic groupings based on DiffServ information. Default is 3.
dot1p	Specifies the priority of traffic groupings based on dot1p information. Default is 2.
priority	Specifies a priority value in the range of 0-15.

Default

access-list = 11
dest-mac = 8
source-mac = 7
diffserv = 3
dot1p = 2
vlan = 1

Usage Guidelines

QoS types with a greater value take higher precedence.

Port-based QoS traffic groupings are always the lowest priority. The priority of port-based traffic cannot be changed.

Example

The following command forces FDB source-mac QoS to take a higher precedence over FDB dest-mac QoS (with a default priority of 8):

```
config qostype priority source-mac 9
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config red drop-probability

```
config red drop-probability <percent>
```

Description

Configures the Random Early Detect (RED) drop-probability.

Syntax Description

percent	Specifies the RED drop probability as a percentage. Range is 0 -100.

Default

N / A.

Usage Guidelines

When the switch detects that traffic is filling up in any of the eight hardware queues, it performs a random discard on subsequent packets, based on the configured RED drop-probability. The destination node detects the dropped packet, and, using standard TCP windowing mechanisms, slows the transmission from the source node. RED drop-probability is configured on a system-wide basis. Only switches and modules with the "*i*" chipset can use RED.

The percentage range is 0 - 100%.

Example

The following command configures the RED drop-probability as 80%:

```
config red drop-probability 80
```

History

This command was first available in ExtremeWare 6.0.10.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan priority

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

Description

Configures the 802.1p priority value for 802.1Q VLAN tags.

Syntax Description

name	Specifies a VLAN name.
priority	Specifies the 802.1p priority value. The value is an integer between 0 and 7.

Default

N/A.

Usage Guidelines

The 802.1p priority field is placed in the 802.1Q tag when a packet is transmitted by the switch.

Example

The following command configures VLAN *accounting* to use priority 6 in its 802.1Q tag:

```
config vlan accounting priority 6
```

History

This command was available in ExtremeWare 6.0

Platform Availability

This command is available on the "*i*" series platforms.

config vlan qosprofile

```
config vlan <name> qosprofile <qosprofile>
```

Description

Configures a VLAN to use a particular QoS profile.

Syntax Description

name	Specifies a VLAN name.
qosprofile	Specifies a QoS profile.

Default

N / A.

Usage Guidelines

"i" series switches support eight QoS profiles (QP1 - QP8). Summit chipset-based switches support four profiles (QP1 - QP4).

Example

The following command configures VLAN *accounting* to use QoS profile QP3:

```
config vlan accounting qosprofile QP3
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

create qosprofile

```
create qosprofile <qosprofile>
```

Description

Creates a QoS profile.

Syntax Description

qosprofile	Specifies a QoS profile name.

Default

N/A.

Usage Guidelines

This command is not available in ExtremeWare v6.0 or later.

Using this command, a new QoS profile is created with the following default values:

- Minimum bandwidth — 0%
- Maximum bandwidth — 100%
- Priority — low

These parameters can then be modified using the `config qosprofile` command.

A QoS profile does not alter the behavior of the switch until it is assigned to a traffic grouping to form a QoS Policy.

Example

The following command creates a QoS profile named *qp4*:

```
create qosprofile qp4
```

History

This command was available in ExtremeWare 2.0.

Support for this command was discontinued in ExtremeWare 6.0.

Platform Availability

This command is available only on platforms based on the Summit chipset.

delete qosprofile

```
delete qosprofile <qosprofile>
```

Description

Deletes a QoS profile created in egress mode.

Syntax Description

qosprofile	Specifies a QoS profile name.

Default

N / A.

Usage Guidelines

This command is not available in ExtremeWare v6.0 or later.

The four default QoS profiles cannot be deleted.

When a QoS profile is removed, all entries previously associated with the QoS profile are changed to use the settings of the default QoS profile named *qp2*.

Example

The following command deletes a QoS profile named *qp1*:

```
delete qosprofile qp1
```

History

This command was available in ExtremeWare 2.0.

Support for this command was discontinued in ExtremeWare 6.0.

Platform Availability

This command is available only on platforms based on the Summit chipset.

disable diffserv examination ports

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

Description

Disables the examination of the Diffserv field in an IP packet.

Syntax Description

portlist	Specifies a list of ports or slots and ports to which the parameters apply. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
all	Specifies that Diffserv examination should be disabled for all ports.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables Diffserv examination on selected ports:

```
disable diffserv examination ports 3,5,6
```

History

This command was available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

disable diffserv replacement ports

```
disable diffserv replacement ports [<portlist> | all]
```

Description

Disables the replacement of diffserv code points in packets transmitted by the switch.

Syntax Description

portlist	Specifies a list of ports or slots and ports to which the parameters apply. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
all	Specifies that Diffserv replacement should be disabled for all ports.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables Diffserv replacement on selected ports:

```
disable diffserv replacement ports 3,5,6
```

History

This command was available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

disable dlcs

```
disable dlcs {fast-ethernet-ports | ports [all | <port_number>]}
```

Description

This command disables WINS snooping for ports on this switch.

Syntax Description

fast-ethernet-ports	Specifies that WINS packet snooping should be disabled on all Fast Ethernet ports.
all	All specifies that WINS packet snooping should be disabled on all ports.
port_number	Specifies a port on which WINS packet snooping should be disabled.

Default

Disabled.

Usage Guidelines

Disabling DLCS means that DLCS information for this switch will no longer be available to the ExtremeWare Enterprise Manager Policy System.

Used with no parameters, this command disables WINS packet snooping on all ports on which it was enabled.

Using the port parameter disabled WINS packet snooping only on the specified port.

Example

The following command disables all WINS packet snooping on the switch:

```
disable dlcs
```

History

This command was available in ExtremeWare 6.0.

The command was modified in ExtremeWare 6.1 to support the `fast-ethernet-ports` parameter.

Platform Availability

This command is available on the "*i*" series platforms.

disable dot1p replacement ports

```
disable dot1p replacement ports [<portlist> | all]
```

Description

Disables the ability to overwrite 802.1p priority values for a given set of ports.

Syntax Description

portlist	Specifies a list of ports or slots and ports to which the parameters apply. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
all	Specifies that 892.1p replacement should be disabled for all ports.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables Diffserv replacement on all ports:

```
disable dot1p replacement ports all
```

History

This command was available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

disable isq vlan

```
disable isq vlan <name>
```

Description

Disables Intra-subnet QoS (ISQ) on a VLAN.

Syntax Description

vlan <name>	Specifies a VLAN name.

Default

Disabled.

Usage Guidelines

None.

Example

This command is not available in ExtremeWare v6.0 or later.

The following command disables ISQ on a VLAN names *accounting*:

```
disable isq vlan accounting
```

History

This command was available in ExtremeWare 4.0.

Support for this command was discontinued in ExtremeWare 6.0.

Platform Availability

This command is available only on platforms based on the Summit chipset.

disable qosmonitor

```
disable qosmonitor
```

Description

Disables the QoS monitoring capability.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables QoS monitoring:

```
disable qosmonitor
```

History

This command was available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

disable red ports

```
disable red ports <portlist>
```

Description

Disables Random Early Detection (RED) on the specified ports.

Syntax Description

portlist	Specifies the port number(s). May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables RED on ports 5-7:

```
disable red ports 5-7
```

History

This command was first available in ExtremeWare 6.0.10.

Platform Availability

This command is available on the "*i*" series platforms.

enable diffserv examination ports

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

Description

Enables the Diffserv field of an ingress IP packet to be examined in order to select a QoS profile.

Syntax Description

portlist	Specifies a list of ports or slots and ports to which the parameters apply. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
all	Specifies that Diffserv examination should be enabled for all ports.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables Diffserv examination on selected ports:

```
enable diffserv examination ports 3,5,6
```

History

This command was available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

enable diffserv replacement ports

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

Description

Enables the diffserv code point to be overwritten in packets transmitted by the switch.

Syntax Description

portlist	Specifies a list of ports or slots and ports to which the parameters apply. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
all	Specifies that Diffserv replacement should be enabled for all ports.

Default

Disabled.

Usage Guidelines

Eight user-defined code points can be configured on each port. The 802.1P priority bits (3-bits) are used to select one of the eight code points.

Example

The following command enables Diffserv replacement on selected ports:

```
enable diffserv replacement ports 3,5,6
```

History

This command was available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

enable dlcs

```
enable dlcs {fast-ethernet-ports | ports [all | <port_number>]}
```

Description

This command enables WINS snooping for ports on the switch.

Syntax Description

fast-ethernet-ports	Specifies that WINS packets should be snooped on all Fast Ethernet ports.
all	Specifies that WINS packets should be snooped on all ports.
port_number	Specifies a port on which WINS packets are to be snooped.

Default

Enables snooping on all ports.

Usage Guidelines

DLCS must be enabled to allow usage of DLCS information by the ExtremeWare Enterprise Manager Policy System.

`enable dlcs` used with no parameters is the same as `enable dlcs ports all`.

The `fast-ethernet-ports` parameter is a shortcut to enable DLCS on all gigabit ethernet ports, rather than having to enter each port individually.

Example

The following command enables DLCS snooping on port 4:

```
enable dlcs ports 4
```

Either of the following commands enable DLCS snooping on all ports:

```
enable dlcs
enable dlcs ports all
```

History

This command was available in ExtremeWare 6.0.

The command was modified in ExtremeWare 6.1 to support the `fast-ethernet-ports` parameter.

Platform Availability

This command is available on the "*i*" series platforms.

enable dot1p replacement ports

```
enable dot1p replacement ports [<portlist> | all]
```

Description

Allows the 802.1p priority field to be overwritten on egress according to the QoS profile to 802.1p priority mapping for a given set of ports.

Syntax Description

portlist	Specifies a list of ports or slots and ports. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
all	Specifies that dot1p replacement should be enabled for all ports.

Default

Disabled.

Usage Guidelines

By default, 802.1p priority information is not replaced or manipulated, and the information observed on ingress is preserved when transmitting the packet.

If 802.1p replacement is enabled, the 802.1p priority information that is transmitted is determined by the hardware queue that is used when transmitting the packet. The mapping is described in Table 11 for switches based on the "*i*" chipset and for other Extreme switches. This mapping cannot be changed.

Table 11: Queue to 802.1p Priority Replacement Value

Hardware Queue Summit Chipset	Hardware Queue "*i*" Chipset	802.1p Priority Replacement Value
Q0	Q0	0
	Q1	1
Q1	Q2	2
	Q3	3
Q2	Q4	4
	Q5	5
Q3	Q6	6
	Q7	7

Example

The following command enables dot1p replacement on all ports:

```
enable dot1p replacement ports all
```

History

This command was available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

enable isq vlan

```
enable isq vlan <name>
```

Description

Enables Intra-Subnet QoS (ISQ) on a per-VLAN basis.

Syntax Description

name	Specifies a VLAN name.

Default

Disabled.

Usage Guidelines

This command is not available in ExtremeWare v6.0 or later.

ISQ allows the application of QoS commands to be effective within a subnet (VLAN) instead of only applying the QoS when traversing a routed subnet. The command syntax for all IPQoS commands remains the same; ISQ is simply enabled on a per VLAN basis.

Because ISQ instructs the switch to look at IP addresses within a VLAN, the normal MAC-based learning and refreshing for layer 2 switching is altered for traffic that matches an IPQoS traffic grouping. Instead, learning and refreshing is done based on IP information in the packets. As a result, it is necessary to increase the FDB aging timer above a normal ARP table refresh time to 50 minutes (3,000 seconds). This occurs automatically when ISQ is enabled. ISQ should not be used on VLANs with clients that have statically defined ARP tables.

Example

The following command enables ISQ on a VLAN named *accounting*:

```
enable isq vlan accounting
```

History

This command was available in ExtremeWare 4.0.

Support for this command was discontinued in ExtremeWare 6.0.

Platform Availability

This command is available only on platforms based on the Summit chipset.

enable qosmonitor

```
enable qosmonitor {port <port>}
```

Description

Enables the QoS monitoring capability on the switch.

Syntax Description

port	Specifies a port.

Default

Disabled.

Usage Guidelines

When no port is specified, the QoS monitor automatically samples all the ports and records the sampled results. Error messages are logged to the syslog if the traffic exceeds the parameters of the QoS profile(s).

The QOS monitor is a utility that monitors the hardware queues associated with any port(s). The QOS monitor keeps track of the number of frames and the frames per second that a specific queue is responsible for transmitting on a physical port. Two options are available: a real-time display and a separate option for retrieving information in the background and writing it to the log.

The real-time display scrolls through the given portlist to provide statistics. The particular port being monitored at that time is indicated by an asterisk (*) appearing after the port number in the display.

Monitoring QoS in the background places transmit counter and any “overflow” information into the switch log. The log notification appears if one of the queues experiences an overflow condition since the last time it was sampled. An overflow entry indicates that a queue was over-subscribed at least temporarily, and is useful for determining correct QoS settings and potential over-subscription issues.

Example

The following command enables the QoS monitoring capability on port 4:

```
enable qosmonitor port 4
```

History

This command was available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

enable red ports

```
enable red ports [mgmt | <portlist>]
```

Description

Enables Random Early Detection (RED) on a port.

Syntax Description

mgmt	Specifies the management port.
portlist	Specifies a list of ports or slots and ports. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

Disabled.

Usage Guidelines

Policy-based QoS can be configured to perform per-port Random Early Detection (RED) and drop-probability. Using this capability, the switch detects when traffic is filling up in any of the eight hardware queues, and performs a random discard on subsequent packets, based on the configured RED drop-probability.

Instead of dropping sessions during times when the queue depth is exceeded, RED causes the switch to lower session throughput. The destination node detects the dropped packet, and, using standard TCP windowing mechanisms, slows the transmission from the source node. RED drop-probability is configured on a system-wide basis, and has a valid range from 0% to 100%. Only switches and modules with the "*i*" chipset can use RED.

Example

The following command enables RED on ports 5-7:

```
enable red ports 5-7
```

History

This command was first available in ExtremeWare 6.0.10.

Platform Availability

This command is available on the "*i*" series platforms.

show dlcs

```
show dlcs
```

Description

Displays the status of DLCS (enabled or disabled) and the status of ports that are snooping WINS packets.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays DLCS status and data from the switch:

```
show dlcs
```

It produces output such as the following:

```
DLCS:         Enabled
Ports:        4
```

History

This command was available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

show dot1p

```
show dot1p
```

Description

Displays the 802.1p-to-QoS profile mappings.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the current 802.1p-to-QoS mappings on the switch:

```
show dot1p
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

show ipqos

```
show ipqos {<ip_destination_address>/<mask>}
```

Description

Displays the IP QoS table.

Syntax Description

ip_destination_address	Specifies a destination IP address.
mask	Specifies a subnet mask.

Default

Displays the entire IP QoS table.

Usage Guidelines

This command is not available in ExtremeWare v6.0 or later.

Including an optional destination IP address shows the IP QoS table entries for that address only.

Example

The following command shows the IP Qos table entries related to the specified subnet:

```
show ipqos 10.1.2.3/24
```

History

This command was available in ExtremeWare 2.0.

Support for this command was discontinued in ExtremeWare 6.0.

Platform Availability

This command is available only on platforms based on the Summit chipset.

show ports qosmonitor

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

Description

Displays real-time QoS statistics for egress packets on one or more ports.

Syntax Description

portlist	Specifies a list of ports or slots and ports. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

Shows QoS statistics for all ports.

Usage Guidelines

The real-time display scrolls through the given portlist to provide statistics. You can choose screens for packet count and packets per second. The specific port being monitored is indicated by an asterisk (*) appearing after the port number in the display.

QoS monitor sampling is configured as follows:

- The port is monitored for 20 seconds before the switch moves on to the next port in the list.
- A port is sampled for five seconds before the packets per second (pps) value is displayed on the screen.

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

show qosprofile

```
show qosprofile {<qosprofile>}
```

Description

Displays QoS information on the switch.

Syntax Description

<qosprofile>	Specifies a QoS profile name.

Default

Displays QoS information for all profiles.

Usage Guidelines

Information displayed includes:

- QoS profile name
- Minimum bandwidth
- Maximum bandwidth
- Priority
- A list of all traffic groups to which the QoS profile is applied

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

show qostype priority

```
show qostype priority
```

Description

Displays QoS traffic grouping priority settings.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the QoS traffic grouping priority settings for this switch:

```
show qostype priority
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

unconfig diffserv examination ports

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

Description

Removes the Diffserv examination code point from a port.

Syntax Description

portlist	Specifies a list of ports or slots and ports. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
all	Specifies that Diffserv examination code points should be removed from all ports.

Default

N/A.

Usage Guidelines

None.

Example

The following command removes Diffserv code-point examination from ports 5-8:

```
unconfig diffserv examination ports 5-8
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

unconfig diffserv replacement ports

```
unconfig diffserv replacement ports [<portlist> | all]
```

Description

Removes the diffserv replacement mapping from a port.

Syntax Description

portlist	Specifies a list of ports or slots and ports. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
all	Specifies that Diffserv replacement mapping should be removed from all ports.

Default

N/A.

Usage Guidelines

None.

Example

The following command removes Diffserv replacement from ports 5-8:

```
unconfig diffserv replacement ports 5-8
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

unconfig qostype priority

```
unconfig qostype priority
```

Description

Resets all traffic grouping priority values to their defaults.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Resets the traffic grouping priorities to the following:

access-list = 11
dest-mac = 8
source-mac = 7
diffserv = 3
dot1p = 2
vlan = 1

Example

The following command resets the QoS traffic grouping priorities:

```
unconfig qostype priority
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

8 Access Policy Commands

This chapter describes the following commands:

- Commands for creating and configuring routing access policies
- Commands for creating and configuring IP access lists
- Commands for creating and configuring route maps

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.

IP access lists (also referred to as just Access Lists or ACLs) consist of IP access rules and 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. Using access lists has no impact on switch performance.

Access lists are typically applied to traffic that crosses layer 3 router boundaries, but it is possible to use access lists within a layer 2 VLAN. Products that use the "*i*" chipset are capable of performing this function with no additional configuration. Products that do not use the "*i*" chipset require the enabling of Intra-subnet QoS (ISQ), to perform this function.

Routing access policies are used to control the advertisement or recognition of routing protocols, such as RIP, OSPF, or BGP. 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.

To use routing access policies, follow these steps:

1 Create an access profile.

2 Configure the access profile mode to be of type *permit*, *deny*, or *none (which allows per-entry configuration of the permit/deny attribute).*

3 Add entries to the access profile.

4 Apply the access profile.

Route maps are used to modify or filter routes redistributed between two routing domains. They are also used to modify or filter the routing information exchanged between the domains.

To use route maps, follow these steps:

1 Create a route map.

2 Add entries to the route map.

3 Add statements to the route map entries.

config access-profile add

```
config access-profile <access_profile> add {<seq_number>} {permit | deny}
[ipaddress <ip address> <mask> {exact} | as-path <path-expression> |
bgp-community [internet | no-export | no-advertise | no-export-subconfed |
<as_no:number> | number <community>] | ipxnet <netid> <netid mask> | ipxsap
<sap_type> <service_name> | vlan]
```

Description

Adds an entry to the access profile.

Syntax Description

access_profile	Specifies an access profile name.
seq-number	Specifies 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	Per-entry permit specification. The per-entry attribute only takes effect if the access-profile mode is none. Otherwise, the overall access profile type takes precedence.
deny	Per-entry deny specification. The per-entry attribute only takes effect if the access-profile mode is none. Otherwise, the overall access profile type takes precedence.
ip address/mask	Specifies an IP address and mask as an entry in the profile list.
exact	Specifies that an exact match with address and mask will be performed. Subnets within the address range will not match entry against entry.
path-expression	Specifies a regular expression string to match against the autonomous system path.
internet	Specifies a match against all routes, because all routes belong to the internet community.
no-export	Match against communities with the no-export attribute.
no-advertise	Match against communities with the no-advertise attribute.
no-export-subconfed	Match against communities with the no-export-subconfed attribute.
as_no:number	Match against a BGP community number, specified in as_no:number format.
community	Match against a BGP community number specified as an unsigned 32-bit integer in decimal format.
netid/netid mask	Specifies an IPX netID and mask as an entry in the profile list.
sap_type/service_name	Specifies an IPX SAP service type and service name as an entry in the profile list.
vlan	Specifies a VLAN name as an entry in the profile list (supported only on BlackDiamond 6800 MSM32 running ExtremeWare 4.1)

Default

N / A.

Usage Guidelines

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.

The explicit sequence number and the permit or deny attribute should be specified if the access profile mode is none.

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.

The `as-path` keyword uses a regular expression string to match against the AS path. Regular expression notation can include any of the characters listed in Table 12.

Table 12: Regular Expression Notation

Character	Definition
N	AS number
N_1 - N_2	Range of AS numbers, where N_1 and N_2 are AS numbers and $N_1 < N_2$
[N_x ... N_y]	Group of AS numbers, where N_x and N_y are AS numbers or a range of AS numbers
[^N_x ... N_y]	Any AS numbers other than the ones in the group
.	Matches any number
^	Matches the beginning of the AS path
$	Matches the end of the AS path
–	Matches the beginning or end, or a space
-	Separates the beginning and end of a range of numbers
*	Matches 0 or more instances
+	Matches 1 or more instances
?	Matches 0 or 1 instance
{	Start of AS SET segment in the AS path
}	End of AS SET segment in the AS path
(	Start of a confederation segment in the AS path
)	End of a confederation segment in the AS path

Example

The following command adds an IP subnet address to access profile *nosales*, as the next available entry:

```
config access-profile nosales add ipaddress 10.1.33.0/24
```

The following command configures the access profile AS1 to permit AS paths beginning with AS number 1, followed by any AS number from 2 - 8, and ending with either AS number 11, 13, or 15:

```
config access-profile AS1 add 15 permit as-path “^1 2-8 [11 13 15]$”
```

History

This form of the command was available in ExtremeWare 6.1. Support for IPX NetID and IPX SAP matching was first available in ExtremeWare 6.2.

A limited version of this command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on "*i*" series platforms. A limited form of the command is available on non-"*i*" series platforms.

config access-profile delete

```
config access-profile <access_profile> delete <seq_number>
```

Description

Deletes an access profile entry using the sequence number.

Syntax Description

access_profile	Specifies an access profile name.
seq-number	Specifies 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.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the entry with sequence number 15 from the access profile AS1:

```
config access-profile AS1 delete 15
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

config access-profile mode

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

Description

Configures the access profile mode to permit or deny access, or to require per-entry access control.

Syntax Description

access_profile	Specifies an access profile name.
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.

Default

Permit.

Usage Guidelines

The access list mode determines whether the items in the list are to be permitted access or denied access.

Example

The following command configures the access profile *no_subnet_33* to deny access:

```
config access-profile no_subnet_33 mode deny
```

The following command specifies that the access profile *no_subnet_33* uses per-entry access control:

```
config access-profile no_subnet_33 mode none
```

History

This command was first available in ExtremeWare 4.0.

The per-entry access control was added in ExtremeWare 6.1.

Platform Availability

This command is available on all platforms.

config route-map add goto

```
config route-map <route_map> <seq_number> add goto <new_route_map>
```

Description

Configures a route map `goto` statement to transfer evaluation to another route map.

Syntax Description

route-map	The name of the route map to which this statement should be added.
seq-number	Specifies the sequence number of the entry in the route map to which this statement should be added.
new-route-map	The name of another route map that should be evaluated.

Default

N / A.

Usage Guidelines

A route map `goto` statement is evaluated only after all `match` and `set` statements have been evaluated.

Example

The following command adds a `goto` statement to entry 25 in route map *map1* that causes evaluation control to transfer to route map *map2*:

```
config route-map map1 25 add goto map2
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the "*i*" series platforms.

config route-map delete goto

```
config route-map <route_map> <seq_number> delete goto <new_route_map>
```

Description

Deletes a route map `goto` statement.

Syntax Description

route-map	The name of the route map from which this statement should be deleted.
seq-number	The sequence number of the entry in the route map from which this statement should be deleted.
new-route-map	The name of another route map.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the `goto` statement from entry 25 in route map *map1* that specifies transfer to route map *map2*:

```
config route-map map1 25 delete goto map2
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the "*i*" series platforms.

config route-map delete

```
config route-map <route_map> delete <seq_number>
```

Description

Deletes an entry from the route map.

Syntax Description

route-map	The name of the route map to which this entry should be added.
seq-number	Specifies a sequence number that uniquely identifies the entry, and determines the position of the entry in the route map.

Default

N / A.

Usage Guidelines

None.

Example

The following command removes the entry with sequence number 20 from the route-map named *bgp-out*:

```
config route-map bgp-out delete 20
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the "*i*" series platforms.

config route-map add match

```
config route-map <route-map> <seq_number> add match [nlri-list
<nlri_access_profile> | as-path [access-profile <as_access_profile> |
<as_number>] | community [access-profile <com_access_profile> |
<as_number>:<number> | number <community> | no-advertise | no-export |
no-export-subconfed] | next-hop <ip address> | med <number> | tag <number>
| origin [igp | egp | incomplete]]
```

Description

Configures a route map match statement.

Syntax Descriptio

route-map	The name of the route map to which this statement should be added.
seq-number	Specifies the sequence number of the entry in the route map to which this statement should be added.
nrli_access_profile	Specifies an access profile against which the NLRI should be matched.
as_access_profile	Specifies an access profile against which the AS path in the path attributes should be matched.
as-number	Specifies an AS number against which the AS path in the path attributes should be matched.
com_access_profile	Specifies a BGP community access profile against which the community attribute should be matched.
as_number:number	Specifies a BGP community number, specified in as_number:number format, against which the community attribute should be matched.
community	Specifies a BGP community number, specified as an unsigned 32-bit integer in decimal format, against which the community attribute should be matched.
no-export	Specifies that the community attribute should be matched against the no-export attribute.
no-advertise	Specifies that the community attribute should be matched against the no-advertise attribute.
no-export-subconfed	Specifies that the community attribute should be matched against the no-export-subconfed attribute.
ipaddress	Specifies an IP address against which the next hop attribute in the path attribute should be matched.
med_number	Specifies a MED number against which the MED in the path attribute should be matched.
origin [igp | egp | incomplete]	Specifies an origin against which the origin in the path attribute should be matched. Values are igp, egp, or incomplete.
tag_number	Specifies a tag value against which the tag associated with the redistributed OSPF route should be matched.

Default

N / A.

Usage Guidelines

A match operation specifies a criteria that must be matched in order for the route to be successful. If there are multiple statements in a route table entry, match statements are evaluated before *set* or *goto* statements.

When an entry has multiple match statements, the primitive `match-one` or `match-all` in the entry determines how many matches are required for success. If an entry has no match statements, the entry is always considered a successful match.

Example

The following command adds a statement to entry 10 in route map *bgp-out* that matches the NLRI against the access profile named *iplist*:

```
config bgp-out 10 add match nlri-list iplist
```

The following command adds a statement to entry 15 in route map *bgp-out* that matches the AS path attribute against the access profile named *aslist*:

```
config bgp-out 15 add match as-path access-profile aslist
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the "*i*" series platforms.

config route-map add set

```
config route-map <route-map> <seq_number> add set [as-path <as_number> |
community [[access-profile <access-profile> | <as_number>:<number> | number
<community> | no-advertise | no-export | no-export-subconfed] | remove |
[add | delete] [access-profile <access-profile> | <as no> : <number> |
number <community> | no-advertise | no-export | no-export-subconfed]] |
next-hop <ip address> | med [internal | <med_number> | remove | [add |
delete] <med_number>] local-preference <number> | weight <number> | origin
[igp | egp | incomplete] | tag <tag_number> | accounting index
<index_number> value <value_number> | cost <number> | cost-type [ase-type-1
| ase-type-2]]
```

Description

Configures a route map set entry.

Syntax Description

route-map	The name of the route map to which this statement should be added.
seq-number	Specifies the sequence number of the entry in the route map to which this statement should be added.
as-number	Prepends the specified AS number to the AS path in the path attribute.
as_access_profile	Sets the community in path attribute to the specified access profile.
as_number:number	Sets the community in path attribute to the specified BGP community number, specified in as_number:number format, in the path attribute.
community	Sets the community in path attribute to the specified BGP community number, specified as an unsigned 32-bit integer in decimal format.
no-export	Sets the community in path attribute to the no-export attribute.
no-advertise	Sets the community in path attribute to the no-advertise attribute.
no-export-subconfed	Sets the community in path attribute to the no-export-subconfed attribute.
remove	Removes the community attribute, if present.
add | delete <as_access_profile>	Adds or deletes the specified access profile to or from the existing community in the path attribute.
add | delete <as_number:number>	Adds or deletes the specified BGP community number, specified in as_number:number format, to or from the existing community in the path attribute.
add | delete <community>	Adds or deletes the specified BGP community number, specified as an unsigned 32-bit integer in decimal format, to or from the existing community in the path attribute.
add | delete <no-export>	Adds or deletes the no-export attribute to or from the existing community in the path attribute.
add | delete <no-advertise>	Adds or deletes he no-advertise attribute to or from the existing community in the path attribute.
add | delete <no-export-subconfed>	Adds or deletes the no-export-subconfed attribute to or from the existing community in the path attribute.
next-hop <ipaddress>	Sets the next hop in the path attribute to the specified IP address.
internal	When used in the BGP neighbor output route map, sets the MED attribute to a value equal to the metric to reach the nexthop.
med_number	Sets the MED attribute to the specified value.

remove	Removes the MED attribute, if present.
add | delete <med_number>	Adds or deletes the specified value to or from the MED that is received. The final result is bound by 0 and 2147483647.
local-preference <number>	Sets the local preference in the path attribute to the specified local preference number.
weight <number>	Sets the weight associated with the NLRI to the specified number.
origin [igp | egp | incomplete]	Sets the origin in the path attributes to the specified origin.
tag <tag_number>	Sets the tag in the route to the specified number.
accounting index <index_number>	Specifies the index number of an accounting index to be set.
value <value_number>	Specifies the value to which the accounting index should be set.
cost <number>	Sets the cost of the route to the specified number.
cost-type <number>	Sets the cost type associated with the route (ase-type-1 or ase-type-2).

Default

N / A.

Usage Guidelines

Route map `set` statements are evaluated after `match` statements, but before the `goto` statement.

Changes to the route maps used to modify or filter NLRI information exchanged with neighbors is immediately effective on the routing information exchanged after the policy changes. The changes can be applied on the NLRI information that had been exchanged before the policy changes by issuing a soft reset on the ingress or egress side, depending on the changes. For soft resets to be applied on the ingress side, the changes must be previously enabled on the neighbor.

Changes to the route maps associated with network aggregation or redistribution commands becomes effective after a maximum interval of 30 seconds. You can immediately apply them by using the soft reconfiguration command.

Example

The following command modify the routing information for a route that matches a statement in entry 15 of route table *bgp-out* include a MED value of 200:

```
config bgp-out 15 add set med 200
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the "*i*" series platforms.

config route-map delete match

```
config route-map <route-map> <seq_number> delete match [nlri-list
<access-profile> | as-path [access-profile <access-profile> | <as_number>]
| community [access-profile <access-profile> | <as_number>:<number> |
number <community> | no-advertise | no-export | no-export-subconfed] |
next-hop <ip address> | med <number> | tag <number> | origin [igp | egp |
incomplete]]
```

Description

Deletes a route map match statement.

Syntax Description

route-map	The name of the route map from which this statement should be deleted.
seq-number	The sequence number of the entry in the route map from which this statement should be deleted.
nrli_access_profile	Specifies an NRLI-list access profile.
as_access_profile	Specifies an AS path access profile.
as-number	Specifies an AS number.
com_access_profile	Specifies a BGP community access profile.
as_number:number	Specifies a BGP community number in as_number:number format.
community	Specifies a BGP community number, specified as an unsigned 32-bit integer in decimal format.
no-export	Specifies the no-export community attribute.
no-advertise	Specifies the no-advertise community attribute.
no-export-subconfed	Specifies the no-export-subconfed community attribute.
ipaddress	Specifies an IP address of the next hop attribute.
med_number	Specifies a MED number.
origin [igp | egp | incomplete]	Specifies an origin. Values are igp, egp, or incomplete.
tag_number	Specifies a tag value.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the statement from entry 15 in route map *bgp-out* that specifies that the access profile *aslist* should be used to match the AS path:

```
config bgp-out 15 add match as-path access-profile aslist
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the "*i*" series platforms.

config route-map delete set

```
config route-map <route-map> <seq_number> delete set [as-path <as_number> |
community [[access-profile <access-profile> | <as_number>:<number> | number
<community> | no-advertise | no-export | no-export-subconfed] | remove |
[add | delete] [access-profile <access-profile> | <as_number>:<number> |
number <community> | no-advertise | no-export | no-export-subconfed]] |
next-hop <ip address> | med <number> | local-preference <number> | weight
<number> | origin [igp | egp | incomplete] | tag <number> | accounting
index <number> value <number> | cost <number> | cost-type [ase-type-1 |
ase-type-2]]
```

Description

Deletes a route map set entry.

Syntax Description

route-map	The name of the route map from which this statement should be deleted.
seq-number	The sequence number of the entry in the route map from which this statement should be deleted.
as-number	Specifies an AS number.
as_access_profile	Specifies an AS path access profile.
as_number:number	Specifies a BGP community number, in as_number:number format.
community	Specifies a BGP community number, as an unsigned 32-bit integer in decimal format.
no-export	Specifies the no-export attribute.
no-advertise	Specifies the no-advertise attribute.
no-export-subconfed	Specifies the no-export-subconfed attribute.
remove	Specifies removing the community attribute.
add | delete <as_access_profile>	Specifies add or delete of the specified access profile.
add | delete <as_number:number>	Specifies add or delete of the specified BGP community number, in as_number:number format.
add | delete <community>	Specifies add or delete of the specified BGP community number, specified as an unsigned 32-bit integer in decimal format.
add | delete <no-export>	Specifies add or delete of the no-export attribute.
add | delete <no-advertise>	Specifies add or delete of the no-advertise attribute.
add | delete <no-export-subconfed>	Specifies add or delete of the no-export-subconfed attribute.
next-hop <ipaddress>	Specifies the IP address of the next hop.
internal	Specifies setting the MED attribute to a value equal to the metric to reach the nexthop.
med_number	Specifies setting the MED attribute to a specified value.
remove	Specifies removing the MED attribute.
add | delete <med_number>	Specifies add or delete of the specified value to or from the MED.
local-preference <number>	Specifies a local preference number.
weight <number>	Specifies a weight associated with the NLRI.

origin [igp | egp | incomplete]	Specifies the origin.
tag <tag_number>	Specifies the tag in the route to the specified number.
accounting index <index_number>	Specifies the index number of an accounting index to be set.
value <value_number>	Specifies a value for the accounting index.
cost <number>	Specifies the cost of the route.
cost-type <number>	Specifies the cost type.

Default

N / A.

Usage Guidelines

None.

Example

The following command deletes the set statement from entry 15 of route table *bgp-out* that specified setting a MED value of 200:

```
config bgp-out 15 delete set med 200
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the "*i*" series platforms.

config route-map add

```
config route-map <route-map> add <seq_number> [permit | deny] {match-one |
match-all}
```

Description

Adds an entry in the route map with the specified sequence number and action.

Syntax Description

route-map	The name of the route map to which this entry should be added.
seq-number	Specifies a sequence number that uniquely identifies the entry, and determines the position of the entry in the route map.
permit	Permits the route.
deny	Denies the route. This is applied only if the match is successful.
match-one	The route map is successful as long as at least one of the matching statements is true.
match-all	The route map is successful only when all match statements are true. This is the default setting.

Default

N/A.

Usage Guidelines

The sequence number determines the order of the entry in the route map.

The action (permit or deny) specifies the action to be taken on a successful match against the statements in the route map.

After an entry has been added to the route map, statements must be added to define the routes that should be matched, using the `config <route-map> add match` command.

Example

The following command adds an entry to the route-map named *bgp-out* that denies all matching routes:

```
config route-map bgp-out add 10 deny
```

The following command adds an entry to the route-map named *bgp-out* that will be evaluated after the previous entry, and that permits all matching routes:

```
config route-map bgp-out add 20 permit
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan access-profile

```
config vlan <name> access-profile [<access_profile> | none]
```

Description

Configures a BlackDiamond 6800 running ExtremeWare 4.1 to control the routing of traffic between VLANs.

Syntax Description

name	Specifies the name of an egress VLAN.
access_profile	Specifies an access profile that contains a list of ingress VLANs.
none	Specifies that no access profile should be associated with this VLAN.

Default

N / A.

Usage Guidelines

This command configures a BlackDiamond 6800 to permit or deny the routing of IP traffic from the specified list of ingress VLANs to the specified egress VLAN. If the access profile uses permit mode, only traffic from the VLANs specified in the access profile will be routed to egress VLANs configured to use that access profile.

The VLAN must already exist. The access profile must be of type VLAN (supported only in ExtremeWare releases 4.0 and earlier).

Example

Given an access profile created and configured as follows:

```
create access-profile okprofile vlan
config access-profile okprofile mode permit
config access-profile okprofile add vlan exec
```

The following command permits traffic from VLAN *exec* to be routed to VLAN *vlan1*:

```
config vlan vlan1 access-profile okprofile
```

History

This command was available in ExtremeWare 4.1.

Support for this command was discontinued in ExtremeWare 6.0.

Platform Availability

This command is available on the BlackDiamond 6800 MSM32 only.

create access-list icmp destination source

```
create access-list <name> icmp destination [<dest_ipaddress>/<mask> | any]
source [<src_ipaddress>/<source_mask> | any] type <icmp_type> code
<icmp_code> [permit | deny] {<portlist>} {precedence <number>}
```

Description

Creates a named IP access list that applies to ICMP traffic.

Syntax Description

name	Specifies the access list name. The access list name can be between 1 and 31 characters.
dest_ipaddress/mask	Specifies an IP destination address and subnet mask. A mask length of 32 indicates a host entry.
	any specifies that any address will match.
src_ipaddress/source_mask	Specifies a source IP address and subnet mask.
	any specifies that any address will match.
icmp_type	Specifies the ICMP_TYPE number. The ICMP type is a number from 0 to 255.
icmp_code	Specifies the ICMP_CODE number. The ICMP code is a number from 0 to 255.
permit	Specifies that packets that match the access list description are permitted to be forward by this switch.
deny	Specifies that packets that match the access list description are filtered (dropped) by the switch.
portlist	Specifies the ingress port(s) on which this rule is applied.
number	Specifies the access list precedence number. The range is 1 to 25,600.

Default

N/A.

Usage Guidelines

The access list is applied to all ingress packets.

Example

This command creates an access list named *denyping* that filters out ping (ICMP echo) packets. ICMP echo packets are defined as type 8 code 0:

```
create access-list denyping icmp destination any source any type 8 code 0 deny ports
any
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

create access-list ip destination source ports

```
create access-list <name> ip destination [<dest_ipaddress>/<mask> | any]
source [<src_ipaddress>/<src_mask> | any] [permit {<qosprofile>} | deny]
ports [<portlist> | any] {precedence <prec_number>}
```

Description

Creates a named IP access list that applies to all IP traffic.

Syntax Description

name	Specifies the access list name. The access list name can be between 1 and 31 characters.
dest_ipaddress/mask	Specifies an IP destination address and subnet mask. A mask length of 32 indicates a host entry.
	any specifies that any address will match.
src_ipaddress/src_mask	Specifies a source IP address and subnet mask.
	any specifies that any address will match.
permit	Specifies that packets that match the access list description are permitted to be forward by this switch.
qosprofile	Specifies an optional QoS profile can be assigned to the access list, so that the switch can prioritize packets accordingly.
deny	Specifies that packets that match the access list description are filtered (dropped) by the switch.
portlist	Specifies the ingress port(s) on which this rule is applied.
	any specifies that the rule will be applied to all ports.
prec_number	Specifies the access list precedence number. The range is 1 to 25,600.

Default

N / A.

Usage Guidelines

The access list is applied to all ingress packets.

Example

The following example defines an access list entry *allow102* with precedence 40 that permits all traffic on any ingress ports to the `10.2.x.x` subnet, and assigns QoS profile Qp3 to those packets:

```
create access-list allow102 ip dest 10.2.0.0/16 source 0.0.0.0/0 permit qosprofile qp3
ports any precedence 40
```

The following command defines a default entry that is used to specify an explicit deny:

```
create access-list denyall ip dest 0.0.0.0/0 source 0.0.0.0/0 deny ports any
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

create access-list tcp destination source ports

```
create access-list <name> tcp destination [<dest_ipaddress>/<mask> | any]
ip-port [<dst_port> | range <dst_port_min> <dst_port_max> | any]
source [<src_ipaddress>/<src_mask> | any] ip-port [<src_port> | range
<src_port_min> <src_port_max> | any] [permit <qosprofile> |
permit-established | deny] ports [<portlist> | any] {precedence
<precedence_num>}
```

Description

Creates a named IP access list that applies to TCP traffic.

Syntax Description

name	Specifies the access list name. The access list name can be between 1 and 31 characters.
dest_ipaddress/mask	Specifies an IP destination address and subnet mask. A mask length of 32 indicates a host entry.
	any specifies that any address will match.
dst_port	Specifies a TCP layer 4 port.
	any specifies that all TCP ports will match.
dst_port_min	Specifies the beginning of a TCP layer 4 port range.
dst_port_max	Specifies the end of a TCP layer 4 port range.
src_ipaddress/src_mask	Specifies a source IP address and subnet mask.
	any specifies that any address will match.
src_port	Specifies a TCP layer 4 port.
	any specifies that all TCP ports will match.
src_port_min	Specifies the beginning of a TCP layer 4 port range.
src_port_max	Specifies the end of a TCP layer 4 port range.
permit	Specifies that packets that match the access list description are permitted to be forward by this switch.
qosprofile	Specifies an optional QoS profile can be assigned to the access list, so that the switch can prioritize packets accordingly.
permit-established	Specifies that a currently-established TCP session is allowed, but TCP packets from source to destination (uni-directional) with SYN=1 and ACK=0 (to initiate a new session) will be dropped.
deny	Specifies that packets that match the access list description are filtered (dropped) by the switch.
portlist	Specifies the ingress port(s) on which this rule is applied.
	any specifies that the rule will be applied to all ports.
prec_number	Specifies the access list precedence number. The range is 1 to 25,600.

Default

N / A.

Usage Guidelines

The access list is applied to all ingress packets.

Example

The following command defines an access-list rule named *allow10_23* with precedence 30 that permits TCP port 23 traffic destined for other 10.x.x.x networks, and assigns QoS profile *Qp4*:

```
create access-list allow10_23 tcp dest 10.0.0.0/8 ip-port 23 source any ip-port any
permit qosprofile qp4 ports any precedence 30
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

create access-list udp destination source ports

```
create access-list <name> udp destination [<dest_ipaddress>/<mask> | any]
ip-port [<dst_port> | range <dst_port_min> <dst_port_max> | any]
source [<src_ipaddress>/<src_mask> | any] ip-port [<src_port> | range
<src_port_min> <src_port_max> | any] [permit <qosprofile> | deny] ports
[<portlist> | any] {precedence <prec_number>}
```

Description

Creates a named IP access list that applies to UDP traffic.

Syntax Description

name	Specifies the access list name. The access list name can be between 1 and 31 characters.
dest_ipaddress/mask	Specifies an IP destination address and subnet mask. A mask length of 32 indicates a host entry.
	any specifies that any address will match.
dst_port	Specifies a UDP layer 4 port.
	any specifies that all UDP ports will match.
dst_port_min	Specifies the beginning of a UDP layer 4 port range.
dst_port_max	Specifies the end of a UDP layer 4 port range.
src_ipaddress/src_mask	Specifies a source IP address and subnet mask.
	any specifies that any address will match.
src_port	Specifies a UDP layer 4 port.
	any specifies that all UDP ports will match.
src_port_min	Specifies the beginning of a UDP layer 4 port range.
src_port_max	Specifies the end of a UDP layer 4 port range.
permit	Specifies that packets that match the access list description are permitted to be forward by this switch.
qosprofile	Specifies an optional QoS profile can be assigned to the access list, so that the switch can prioritize packets accordingly.
deny	Specifies that packets that match the access list description are filtered (dropped) by the switch.
portlist	Specifies the ingress port(s) on which this rule is applied.
	any specifies that the rule will be applied to all ports.
prec_number	Specifies the access list precedence number. The range is 1 to 25,600.

Default

N/A.

Usage Guidelines

The access list is applied to all ingress packets.

Example

The following command defines an access-list rule named *allow10_35* with precedence 70 that permits udp port 35 traffic destined for other 10.X.X.X networks, and assigns QoS profile *Qp2*:

```
create access-list allow10_35 udp dest 10.0.0.0/8 ip-port 35 source any ip-port any
permit qosprofile qp2 ports any precedence 70
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

create access-profile

```
create access-profile <access_profile> type [ipaddress | ipx-node | ipx-net
| ipx-sap | as-path | bgp-community | vlan]
```

Description

Creates an access profile.

Syntax Description

access_profile	Specifies an access profile name.
ipaddress	Specifies that the profile entries will be a list of IP address/mask pairs.
ipx-node	Specifies that the profile entries will be a list of IPX node addresses.
ipx-net	Specifies that the profile entries will be a list of IPX NetIDs.
ipx-sap	Specifies that the profile entries will be a list of IPX SAP advertisements.
as-path	Specifies that the profile entries will be a list of AS path expressions.
bgp-community	Specifies that the profile entries will be a list of BGP community numbers.
vlan	Specifies that the profile entries will be a list of VLANs (supported only on BlackDiamond 6800 MSM32 running ExtremeWare 4.1)

Default

N/A.

Usage Guidelines

You must give the access profile a unique name (in the same manner as naming a VLAN, protocol filter, or Spanning Tree Domain).

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.

For version 4.0:

- Only type `ipaddress` was supported, and the `type` keyword was not used.
- On BlackDiamond 6800 MSM32 running ExtremeWare 4.1, the VLAN keyword specifies that profile entries will be a list of VLANs.

Example

The following command creates an access profile named *nosales* that will contain IP address/mask pairs:

```
create access-profile nosales type ipaddress
```

The following command creates an access profile that will contain AS path expressions:

```
create access-profile AS1 type as-path
```

History

This form of the command was available in ExtremeWare 6.1. Support for the IPX node, NetID and SAP advertisement types was added in ExtremeWare 6.2.

A limited version of this command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on "*i*" series platforms. A limited form of the command is available on platforms based on the Summit chipset.

create route-map

```
create route-map <route_map>
```

Description

Creates a route map statement.

Syntax Description

route_map	Specifies a route map name.

Default

N / A.

Usage Guidelines

Route maps are a mechanism that can be used to conditionally control the redistribution of routes between two routing domains, and to modify the routing information that is redistributed.

Route maps are used in conjunction with the match and set operations. A match operation specifies a criteria that must be matched. A set operation specifies a change that is made to the route when the match operation is successful.

After a route map statement has been created, you must add entries to the route-map, and then add statements to the route map entries.

Example

The following command creates a route-map named *bgp-out*:

```
create route-map bgp-out
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the "*i*" series platforms.

delete access-list

```
delete access-list [<name> | all]
```

Description

Deletes an access list.

Syntax Description

name	Specifies the name of the access list to be deleted.
all	Specifies that all access lists should be deleted.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes access list *allow102*:

```
delete access-list allow102
```

History

This command was first available in ExtremeWare 6.0.

The command was modified in ExtremeWare 6.2.1 to provide the `all` option.

Platform Availability

This command is available on all *"i"* series platforms.

delete access-profile

```
delete access-profile <access_profile>
```

Description

Deletes an access profile.

Syntax Description

access_profile	Specifies an access profile name.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes an access profile named *nosales*:

```
delete access-profile nosales
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

delete route-map

```
delete route-map <route_map>
```

Description

Deletes a route map statement from the route map.

Syntax Description

route_map	Specifies a route map name.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes a route-map named *bgp-out*:

```
delete route-map bgp-out
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the "*i*" series platforms.

disable access-list

```
disable access-list <name> [counter | log]
```

Description

Disables message logging or the collection of access-list statistics.

Syntax Description

name	Specifies the name of the access list.
counter	Specifies that access-list statistics collection should be disable.
log	Specifies that message logging to the Syslog facility for each packet that matches the access-list description should be disabled.

Default

Counting is ON, logging is OFF.

Usage Guidelines

None.

Example

The following command disables statistics collection for access list *allow102*:

```
disable access-list allow102 counter
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

enable access-list

```
enable access-list <name> [counter | log]
```

Description

Enables message logging or the collection of access-list statistics.

Syntax Description

name	Specifies the name of the access list.
counter	Specifies that access-list statistics should be collected.
log	Specifies that a message should be logged to the Syslog facility for each packet that matches the access-list description. The message details the properties of the packet.

Default

Counting is ON, logging is OFF.

Usage Guidelines

None.

Example

The following command enables statistics collection for access list *allow102*:

```
enable access-list allow102 counter
```

The following command enables logging of packets for access list *allow102*:

```
enable access-list allow102 log
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

show access-list

```
show access-list {<name> | port <portlist>}
```

Description

Displays access list information and real-time statistics.

Syntax Description

name	Specifies the name of an access list to be displayed.
portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

Shows information for all access lists.

Usage Guidelines

To verify access list settings, you can view the access list configuration and see real-time statistics on which access list entries are being accessed when processing traffic.

Example

The following command shows information on all current the access lists:

```
show access-list
```

It produces output similar to the following:

```
Rule         Dest/mask:L4DP         Src/mask:L4SP        Flags  Hits
test1        0.0.0.0/ 0:   0          0.0.0.0/ 0:   0     I-P-X   1531

Flags: I=IP, T=TCP, U=UDP, E=Established, M=ICMP
       P=Permit Rule, D=Deny Rule
       N=Port Specific Rule, X=Any Port
```

The following command shows real-time access list statistics for ingress ports 5-7:

```
show access-list port 5-7
```

The following command shows information for access list *test1*:

```
show access-list test1
```

The command generates output similar to the following:

```
test1
       Protocol: ip    Action: permit qp1
       Destination: 0.0.0.0/0  any
       Source:  any    any
       Precedence: 0
       Rule Number: 0
       Hit Count: 4566 Flags: ac
       Ports:
              any
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

show access-list-fdb

```
show access-list-fdb
```

Description

Displays the hardware access control list mapping.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the hardware access control list mapping:

```
show access-list-fdb
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

show access-list-monitor

```
show access-list-monitor
```

Description

Initiates the access-list information display, and refreshes it until discontinued.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command initiates a display of real-time access list information. Use the keys as shown in Table 13 to change the view of the data. The [Esc] or [Return] keys will discontinue the display.

Table 13: 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.

Example

The following command initiates the access-list information display:

```
show access-list-monitor
```

The command displays output similar to the following:

```
Access List     Proto   Destination     Source          Hit Count
=========================================================================
test1           ip      0.0.0.0/0       0.0.0.0/0         1922
```

The Hit Count continues to be updated until you exit from the display or enter “0” to reset the count to zero.

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the “*i*” series platforms.

show access-profile

```
show access-profile {<access_profile>}
```

Description

Displays access-profile related information for the switch.

Syntax Description

access_profile	Specifies an access profile.

Default

Shows all access profile information for the switch.

Usage Guidelines

None.

Example

The following command displays access-profile related information for access profile *nosales*:

```
show access-profile nosales
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

9 NAT Commands

This chapter covers the following topics:

- Configuring VLANs for Network Address Translation (NAT)
- Configuring NAT translation rules
- Displaying NAT settings

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 (any Extreme Networks switch using the "*i*" chipset) rewrite the source IP address and layer 4 port of the packets.

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 using rules that specify the IP subnets involved and the algorithms used to translate the addresses.

NOTE

The NAT modes in ExtremeWare 6.2 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.

TCP and UDP 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.

config nat add vlan map

```
config nat add vlan <name> map source [any | <source_ipaddress>/<mask>]
{l4-port [any | <port> {- <port>}]}
{destination <dest_ipaddress>/<mask> {l4-port [any | <port> {- <port>}]}}
to <ip address> [/<mask> | - <ip address>]
[tcp | udp | both] [portmap {<min> - <max>} | auto-constrain]
```

Description

Adds a NAT translation rule that translates private IP addresses to public IP addresses on the outside VLAN.

Syntax Description

name	Specifies the name of the outside VLAN to which this rule applies.
source_ipaddress/mask	Specifies a subnet IP address (in the format x.x.x.x/mask where mask is the number of bits in the subnet mask) that defines the source of the traffic to be mapped.
l4-port	Specifies a layer 4 port or port range. When used with a source IP address, indicates that the rule applies only to traffic from the specified layer 4 port(s).
	When used with a destination IP address, indicates that the rule applies only to packets with the specified layer 4 port(s) as their destination.
port	Specifies a port number in the range 1 to 65535.
	any indicates that the rule should be applied to traffic to/from any layer 4 port.
dest_ipaddress/mask	Specifies a subnet IP address (in the format x.x.x.x/mask where mask is the number of bits in the subnet mask) used to determine the packets to which this rule applies.
nat_ipaddress	Specifies an IP address for the outside VLAN to which the source IP addresses will be mapped. This can be specified as a subnet (IP address and mask) or as an address range.
tcp	Specifies only TCP traffic should be translated.
udp	Specifies only UDP traffic should be translated.
both	Specifies that both TCP and UDP traffic should be translated.
portmap	Specifies that port-mapping mode should be used.
min	Specifies a port number in the range 1 to 65535.The default setting is 1024.
max	Specifies a port number in the range 1 to 65535. The default setting is 65535.
auto-constrain	Specifies that each inside IP address should be restricted in the number of simultaneous connections.

Default

N / A.

Usage Guidelines

Four different modes are used to determine how the outside IP addresses and layer 4 ports are assigned:

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

When static mapping is used, each inside IP address uses a single outside IP address. The layer 4 ports are not changed, and only the IP address is rewritten.

With dynamic mapping, 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. The layer 4 ports are not changed. When the last session for a specific inside IP address closes, that outside IP address can be used by other hosts.

The `source` IP address specifies private side IP addresses and the `to` IP address (the NAT address) specifies the public side IP address. The addition of the `destination` optional keyword after the source IP address and mask species that the NAT rule to be applied to only packets with a specific destination IP address.

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.

With static or dynamic translation rules, which do not rely on layer 4 ports, ICMP traffic is translated and allowed to pass.

The addition of a layer 4 protocol name and the `portmap` keyword tells the switch to use portmap mode. 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.

Optionally, you may specify the range of layer 4 ports the switch chooses on the translated IP addresses. The default setting for `min` is 1024. The default setting for `max` is 65535. There is a performance penalty associated with specifying a specific port range other than the default.

ICMP traffic is not translated in portmap mode. You must add a dynamic NAT rule for the same IP address range to allow for ICMP traffic.

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 auto-constrain mode. You must add a dynamic NAT rule for the same IP address range to allow for ICMP traffic.

The addition of the l4-port optional keyword allows the NAT rule to be applied to only packets with a specific layer 4 source or destination port. If you use the layer 4-port command after the source IP/mask, the rule will only match if the port(s) specified are the source layer 4-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 layer 4 ports. Both options may be used together to further limit the rule. If you specify layer 4 ports, ICMP traffic will not translated and allowed to pass.

Rules are processed in order, usually in the order in which they were added. When a single rule is matched, no other rules are processed. You can view the rule order using the show nat rules command.

Example

The following command defines a static translation rule that specifies that traffic coming from 192.168.1.12 be mapped to 216.52.8.32 on outside VLAN out_vlan_1:

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

The following command defines a dynamic translation rule that specifies that traffic coming from subnet 192.168.1.0 should be mapped to IP addresses in the range of 216.52.8.1 to 216.52.8.31 on outside VLAN out_vlan_1:

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

The following command defines a translation rule that specifies that TCP/UDP packets coming from 192.168.1.12 and destined for 192.168.5.20 be mapped to 216.52.8.32 on outside VLAN out_vlan_1:

```
config nat add out_vlan_1 map source 192.168.1.12/32 destination 192.168.5.20 to
216.52.8.32/32
```

The following command defines a portmap translation rule that specifies that both TCP and UDP traffic from subnet 102.168.2.0/25 be mapped to available layer 4 ports on the IP addresses in the subnet 216.52.8.32/28:

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

The following command defines a portmap translation rule that specifies that only TCP traffic from subnet 102.168.2.0/25 be mapped to layer 4 ports in the range of 1024-8192 on the IP addresses in the subnet 216.52.8.32/28:

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

The following command specifies an autoconstrain NAT translation rule that applies to both TCP and UDP traffic:

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

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config nat delete

```
config nat delete [all |
vlan <name> map source [any | <ip address>/<mask>]
  {l4-port [any | <port> {- <port>}]}
{destination <ip address>/<mask> {l4-port [any | <port> {- <port>}]}}
to <ip address> [/<mask> | - <ip address>]
[tcp | udp | both] [portmap {<min> - <max>} | auto-constrain]
```

Description

Deletes a NAT translation rule.

Syntax Description

all	Specifies that all NAT rules should be deleted.
name	Specifies the name of the outside VLAN to which this rule applies.
source_ipaddress/mask	Specifies a subnet IP address (in the format x.x.x.x/mask where mask is the number of bits in the subnet mask) that defines the source of the traffic to be mapped.
l4-port	Specifies a layer 4 port or port range. When used with a source IP address, indicates that the rule applies only to traffic from the specified layer 4 port(s).
	When used with a destination IP address, indicates that the rule applies only to packets with the specified layer 4 port(s) as their destination.
port	Specifies a port number in the range 1 to 65535.
	any indicates that the rule should be applied to traffic to/from any layer 4 port.
dest_ipaddress/mask	Specifies a subnet IP address (in the format x.x.x.x/mask where mask is the number of bits in the subnet mask) used to determine the packets to which this rule applies.
nat_ipaddress	Specifies an IP address for the outside VLAN to which the source IP addresses will be mapped. This can be specified as a subnet (IP address and mask) or as an address range.
tcp	Specifies only TCP traffic should be translated.
udp	Specifies only UDP traffic should be translated.
both	Specifies that both TCP and UDP traffic should be translated.
min	Specifies a port number in the range 1 to 65535. The default setting is 1024.
max	Specifies a port number in the range 1 to 65535. The default setting is 65535.
autoconstrain	Specifies that each inside IP address should be restricted in the number of simultaneous connections.

Default

N / A.

Usage Guidelines

To delete all NAT rules, use the `all` keyword. To delete a specific NAT rule, you must use exactly the same parameters that you used to create the rule.

Example

The following command deletes a portmap translation rule:

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

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config nat finrst-timeout

```
config nat finrst-timeout <seconds>
```

Description

Configures the timeout for a TCP session that has been torn down or reset.

Syntax Description

seconds	Specifies the number of seconds to wait before the session table entry times out.

Default

Default timeout is 60 seconds.

Usage Guidelines

Setting the timeout to zero specifies that session table entries should not be timed-out. This is not normally recommended as NAT resources will get used up.

Example

The following command configures the timeout for a reset or torn-down TCP session to 120 seconds:

```
config nat finrst-timeout 120
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config nat icmp-timeout

```
config nat icmp-timeout <seconds>
```

Description

Configures the timeout for an ICMP packet.

Syntax Description

seconds	Specifies the number of seconds to wait before the session table entry times out.

Default

Default timeout is 3 seconds.

Usage Guidelines

Setting the timeout to zero specifies that session table entries should not be timed-out. This is not normally recommended as NAT resources will get used up.

Example

The following command configures the timeout for an ICMP packet to 5 seconds:

```
config nat icmp-timeout 5
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config nat syn-timeout

```
config nat syn-timeout <seconds>
```

Description

Configures the timeout for an entry with an unacknowledged TCP SYN state.

Syntax Description

seconds	Specifies the number of seconds to wait before the session table entry times out.

Default

Default timeout is 60 seconds.

Usage Guidelines

Setting the timeout to zero specifies that session table entries should not be timed-out. This is not normally recommended as NAT resources will get used up.

Example

The following command configures the timeout for a session with an unacknowledged SYN packet to 120 seconds:

```
config nat syn-timeout 120
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config nat tcp-timeout

```
config nat tcp-timeout <seconds>
```

Description

Configures the timeout for a fully setup TCP SYN session.

Syntax Description

seconds	Specifies the number of seconds to wait before the session table entry times out.

Default

Default timeout is 120 seconds.

Usage Guidelines

Setting the timeout to zero specifies that session table entries should not be timed-out. This is not normally recommended as NAT resources will get used up.

Example

The following command configures the timeout for a TCP session to 90 seconds:

```
config nat tcp-timeout 90
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config nat timeout

```
config nat timeout <seconds>
```

Description

Configures the timeout for any IP packet that is not TCP, UDP, or ICMP.

Syntax Description

seconds	Specifies the number of seconds to wait before the session table entry times out.

Default

Default timeout is 600 seconds.

Usage Guidelines

Setting the timeout to zero specifies that session table entries should not be timed-out. This is not normally recommended as NAT resources will get used up.

Example

The following command configures the timeout for packets other than TCP, UDP, or ICMP to 240 seconds:

```
config nat timeout 240
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config nat udp-timeout

```
config nat udp-timeout <seconds>
```

Description

Configures the timeout for a UDP session.

Syntax Description

seconds	Specifies the number of seconds to wait before the session table entry times out.

Default

Default timeout is 120 seconds.

Usage Guidelines

Setting the timeout to zero specifies that session table entries should not be timed-out. This is not normally recommended as NAT resources will get used up.

Example

The following command configures the timeout for a UDP session to 90 seconds:

```
config nat udp-timeout 90
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config nat vlan

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

Description

Configures a VLAN to participate in NAT.

Syntax Description

vlan_name	Specifies a VLAN name.
inside	Specifies that the VLAN is an inside VLAN.
outside	Specifies that the VLAN is an outside VLAN.
none	Disables NAT functions on this VLAN.

Default

N / A.

Usage Guidelines

When a VLAN is configured to be `inside`, traffic from that VLAN is translated only if it has a matching NAT rule. Any unmatched traffic will be routed normally and not be translated. When a VLAN is configured to be `outside`, it routes all traffic.

Because all traffic runs through the central processing unit (CPU), it cannot run at line-rate.

Normally, outside traffic will be able to initiate connections to the internal private IP addresses. If you want to prevent this, you can create IP and ICMP access-lists on the outside VLAN ports to deny traffic destined for the inside IP addresses. There is a NAT performance penalty when you do this.

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

Example

The following command configures the VLAN out_vlan_1 as an outside VLAN for use with NAT:

```
config nat vlan out_vlan_1 outside
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

disable nat

```
disable nat
```

Description

Disables network address translation on the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables NAT functionality on the switch:

```
disable nat
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

enable nat

```
enable nat
```

Description

Enables network address translation on the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables NAT functionality on the switch:

```
enable nat
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

show nat

```
show nat {timeout | stats | connections | rules {vlan <outside_vlan>}}
```

Description

Displays NAT settings.

Syntax Description

timeout	Specifies the display of NAT timeout settings.
stats	Specifies the display of statistics for NAT traffic.
connections	Specifies the display of the current NAT connection table.
rules	Specifies the display of NAT rules, optionally for a specific VLAN.
outside_vlan	Specifies the outside VLAN for which NAT rules should be displayed.

Default

Displays all NAT settings.

Usage Guidelines

Use the keyword `stats` to display statistics for the NAT traffic, including:

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

Use the keyword `connections` to display the current NAT connection table, including source IP/layer 4 port mappings from inside to outside.

Use the keyword `rules` to display the NAT translation rules for the outside VLANs configured on the switch. Rules are displayed in the order they are processed, starting with the first one. To display the NAT rules for a specific VLAN, add the VLAN name.

Use the keyword `timeout` to display the NAT timeout settings configured on the switch.

Example

The following command shows the NAT translation rules configured for VLAN *out_vlan_1*:

```
show nat rules vlan out_vlan_1
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

10 SLB Commands

This chapter discusses server load balancing (SLB) and flow redirect commands.

SLB transparently distributes client requests among several servers. The main use for SLB is for web hosting (using redundant servers to increase the performance and reliability of busy websites).

You can use SLB to manage and balance traffic for client equipment such as web servers, cache servers, routers, and proxy servers. SLB is especially useful for e-commerce sites, Internet service providers, and managers of large intranets.

SLB also provides health checking. Health checking allows you to actively poll nodes to determine their health. The switch makes new connections only if the virtual server and node are both enabled and passing health checks. The switch considers a virtual server or node active unless a health check fails. If a health check fails, the switch considers the virtual server or node inactive. A virtual server or node is also considered inactive if it is disabled and has zero active connections.

Flow redirect overrides routing decisions to transparently redirect client requests to a target device (or group of devices). Unlike SLB, you do not duplicate content on the target device(s).

The switch can only redirect traffic that crosses a VLAN boundary, because flow redirect operates at layer 3. Flow redirection examines traffic and redirects it based on the following criteria, in order of priority:

1 Destination IP address and mask

2 Layer 4 port

3 Source IP address and mask

You can use flow redirect for the following:

- Web cache redirection
- Policy-based routing

clear slb connections

```
clear slb connections {ipaddress <ip address> : <port> | vip <vipname>}
```

Description

Clears all existing SLB connections.

Syntax Description

ip address	Specifies an IP address.
port	Specifies a port.
vipname	Specifies a virtual server.

Default

N / A.

Usage Guidelines

If you do not specify an IP address or a virtual server, all connections are cleared.

This interrupts all current connections, but does not prevent new connections from being established. To prevent new connections from being established, disable SLB to each virtual server using the following command:

```
disable slb vip <vipname> all
```

To prevent new connections from being established to a specific virtual server and simultaneously close all current connections, use the following command:

```
disable slb vip <vipname> all close-connections-now
```

Example

The following command clears the connections to the virtual server “content”:

```
clear slb connections content
```

History

This command was available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

clear slb vip persistence

```
clear slb vip [<vipname> | all] persistence
```

Description

Clears the connection information in the persistence table.

Syntax Description

vipname	Specifies a virtual server.
all	Specifies all virtual servers.

Default

N/A.

Usage Guidelines

Use this command only during testing. Clearing persistence disables applications, such as shopping carts, that require persistence.

Example

The following command clears all information in the persistence table:

```
clear slb vip all persistence
```

History

This command was available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config flow-redirect add next-hop

```
config flow-redirect <flow redirect> add next-hop <ip address>
```

Description

Adds the next hop host (gateway) that is to receive the packets that match the flow redirect policy.

Syntax Description

flow redirect	Specifies a flow redirect policy.
ip address	Specifies an IP address.

Default

N / A.

Usage Guidelines

This command also automatically enables ping-based health checking.

Example

The following command adds the next hop of 10.2.1.20 to the flow redirect policy named “http_flow”:

```
config flow-redirect http_flow add next-hop 10.2.1.20
```

History

This command was available in ExtremeWare 6.1.4.

Platform Availability

This command is available on all “*i*” series platforms.

config flow-redirect delete next-hop

```
config flow-redirect <flow redirect> delete next-hop <ip address>
```

Description

Deletes the next hop host (gateway).

Syntax Description

flow redirect	Specifies a flow redirect policy.
ip address	Specifies an IP address.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the next hop of 10.2.1.20 from the flow redirect policy named "http_flow":

```
config flow-redirect http_flow delete next-hop 10.2.1.20
```

History

This command was available in ExtremeWare 6.1.4.

Platform Availability

This command is available on all "*i*" series platforms.

config flow-redirect service-check ftp

```
config flow-redirect <flow redirect> service-check ftp user <user name>
<password>
```

Description

Configures the flow redirect FTP check.

Syntax Description

flow redirect	Specifies a flow redirect policy.
user name	Specifies the user name for logging in to the FTP service.
password	Specifies the password for logging in to the FTP service.

Default

N / A.

Usage Guidelines

This command automatically enables FTP check. The FTP check logs in to each next hop specified in the flow redirect policy with the user name and password specified.

For ExtremeWare 6.2.0 and prior, the frequency of the FTP check is 60 seconds, the timeout of the FTP check is 180 seconds, and you cannot configure these times.

For ExtremeWare 6.2.1 and later, configure the frequency and timeout using the following command:

```
config flow-redirect timer service-check
```

Example

The following command configures (and enables) FTP check for the flow redirect policy named “ftp_flow” and logs in with the user name “test” and password “extreme”:

```
config flow-redirect ftp_flow service-check ftp user test extreme
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all “*i*” series platforms.

config flow-redirect service-check http

```
config flow-redirect <flow redirect> service-check http url <url>
match-string <alphanumeric string>
```

Description

Configures the flow redirect HTTP check.

Syntax Description

flow redirect	Specifies a flow redirect policy.
url	Specifies the URL to be checked.
alphanumeric string	Specifies the text to search for.

Default

N/A.

Usage Guidelines

This command automatically enables HTTP check. The HTTP requests the designated URL from each next hop specified in the flow redirect policy and checks for the specified alphanumeric string in the first 5000 bytes. Extreme Networks recommends that you create a specific URL dedicated to this check.

Do not include "http://" in the URL. To check a URL beyond the root directory, include the path in the specified URL. The maximum length of a URL is 255 characters.

For ExtremeWare 6.2.0 and prior, the frequency of the HTTP check is 60 seconds, the timeout of the HTTP check is 180 seconds, and you cannot configure these times.

For ExtremeWare 6.2.1 and later, configure the frequency and timeout using the following command:

```
config flow-redirect timer service-check
```

Example

The following command configures (and enables) HTTP check for the flow redirect policy named "http_flow" and checks http://www.checktest.com for the string "test":

```
config flow-redirect http_flow service-check http url www.checktest.com match-string
test
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all "*i*" series platforms.

config flow-redirect service-check L4-port

```
config flow-redirect <flow redirect> service-check L4-port
```

Description

Configures the flow redirect layer 4 port check.

Syntax Description

flow redirect	Specifies a flow redirect policy.

Default

N / A.

Usage Guidelines

This command automatically enables layer 4 port check. The layer 4 port check opens and closes the layer 4 port specified in the flow redirect policy.

For ExtremeWare 6.2.0 and prior, the frequency of the layer 4 port check is 10 seconds, the timeout of the layer 4 port check is 30 seconds, and you cannot configure these times.

For ExtremeWare 6.2.1 and later, configure the frequency and timeout using the following command:

```
config flow-redirect timer tcp-port-check
```

Example

The following command configures (and enables) layer 4 port check for the flow redirect policy named “http_flow”:

```
config flow-redirect http_flow service-check L4-port
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all “*i*” series platforms.

config flow-redirect service-check nntp

```
config flow-redirect <flow redirect> service-check nntp <newsgroup>
```

Description

Configures the flow redirect NNTP check.

Syntax Description

flow redirect	Specifies a flow redirect policy.
newsgroup	Specifies the news group to be checked.

Default

N/A.

Usage Guidelines

This command automatically enables NNTP check. The NNTP check checks the news server specified in the flow redirect policy.

For ExtremeWare 6.2.0 and prior, the frequency of the NNTP check is 60 seconds, the timeout of the NNTP check is 180 seconds, and you cannot configure these times.

For ExtremeWare 6.2.1 and later, configure the frequency and timeout using the following command:

```
config flow-redirect timer service-check
```

Example

The following command configures (and enables) NNTP check for the flow redirect policy named “nttp_flow” and checks the newsgroup “testgroup”:

```
config flow-redirect nttp_flow service-check nntp testgroup
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all “*i*” series platforms.

config flow-redirect service-check ping

```
config flow-redirect <flow redirect> service-check ping
```

Description

Configures the flow redirect ping check.

Syntax Description

flow redirect	Specifies a flow redirect policy.

Default

N / A.

Usage Guidelines

This command automatically enables ping check.

Ping check is also automatically enabled when you add a next hop using the following command:

`config flow-redirect add next-hop`

In ExtremeWare 6.2.0 and prior, the frequency of the ping check is 10 seconds, the timeout of the ping check is 30 seconds, and you cannot configure these times.

For ExtremeWare 6.2.1 and later, configure the frequency and timeout using the following command:

`config flow-redirect timer ping-check`

Example

The following command configures (and enables) ping check for the flow redirect policy named "http_flow":

`config flow-redirect http_flow service-check ping`

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all "*i*" series platforms.

config flow-redirect service-check pop3

```
config flow-redirect <flow redirect> service-check pop3 user <user name>
<password>
```

Description

Configures the flow redirect POP3 check.

Syntax Description

flow redirect	Specifies a flow redirect policy.
user name	Specifies the user name for logging in to the POP3 service.
password	Specifies the password for logging in to the POP3 service.

Default

N / A.

Usage Guidelines

This command automatically enables POP3 check. The POP3 check logs in to each next hop specified in the flow redirect policy with the user name and password specified.

For ExtremeWare 6.2.0 and prior, the frequency of the POP3 check is 60 seconds, the timeout of the POP3 check is 180 seconds, and you cannot configure these times.

For ExtremeWare 6.2.1 and later, configure the frequency and timeout using the following command:

```
config flow-redirect timer service-check
```

Example

The following command configures (and enables) POP3 check for the flow redirect policy named “pop3_flow” and logs in with the user name “test” and the password “extreme”:

```
config flow-redirect pop3_flow service-check pop3 user test extreme
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all “*i*” series platforms.

config flow-redirect service-check smtp

```
config flow-redirect <flow redirect> service-check smtp <dns domain>
```

Description

Configures the flow redirect SMTP check.

Syntax Description

flow redirect	Specifies a flow redirect policy.
dns domain	Specifies the DNS domain of the mail server.

Default

N / A.

Usage Guidelines

This command automatically enables SMTP check. The SMTP ensures that the mail server specified in the flow redirect policy is able to send and receive mail.

For ExtremeWare 6.2.0 and prior, the frequency of the SMTP check is 60 seconds, the timeout of the SMTP check is 180 seconds, and you cannot configure these times.

For ExtremeWare 6.2.1 and later, configure the frequency and timeout using the following command:

```
config flow-redirect timer service-check
```

Example

The following command configures (and enables) SMTP check for the flow redirect policy named "smtp_flow":

```
config flow-redirect smtp_flow service-check smtp 10.4.1.40
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all "*i*" series platforms.

config flow-redirect service-check telnet

```
config flow-redirect <flow redirect> service-check telnet user <user name>
<password>
```

Description

Configures the flow redirect Telnet check.

Syntax Description

flow redirect	Specifies a flow redirect policy.
user name	Specifies the user name for logging in to the telnet service.
password	Specifies the password for logging in to the telnet service.

Default

N / A.

Usage Guidelines

This command automatically enables Telnet check. The Telnet check logs in to each next hop specified in the flow redirect policy with the user name and password specified.

For ExtremeWare 6.2.0 and prior, the frequency of the Telnet check is 60 seconds, the timeout of the Telnet check is 180 seconds, and you cannot configure these times.

For ExtremeWare 6.2.1 and later, configure the frequency and timeout using the following command:

```
config flow-redirect timer service-check
```

Example

The following command configures (and enables) Telnet check for the flow redirect policy named “telnet_flow” and logs in with the user name “test” and the password “extreme”:

```
config flow-redirect telnet_flow service-check telnet user test extreme
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all “*i*” series platforms.

config flow-redirect timer ping-check

```
config flow-redirect timer ping-check frequency <seconds> timeout <seconds>
```

Description

Configures the flow redirect ping-check frequency and timeout.

Syntax Description

frequency	Specifies the ping-check frequency. The range is 1 to 60.
timeout	Specifies the ping-check timeout. The range is 1 to 60.

Default

The default frequency is 10 seconds.

The default timeout is 30 seconds.

Usage Guidelines

The frequency must be less than the timeout.

Example

The following command configures a flow redirect ping-check frequency of 5 seconds and a timeout of 15 seconds:

```
config flow-redirect timer ping-check frequency 5 timeout 15
```

History

This command was first available in ExtremeWare 6.2.1.

Platform Availability

This command is available on all "*i*" series platforms.

config flow-redirect timer service-check

```
config flow-redirect timer service-check frequency <seconds> timeout
<seconds>
```

Description

Configures the flow redirect service-check frequency and timeout.

Syntax Description

frequency	Specifies the service-check frequency. The range is 15 to 300.
timeout	Specifies the service-check timeout. The range is 15 to 300.

Default

The default frequency is 60 seconds.

The default timeout is 180 seconds.

Usage Guidelines

The frequency must be less than the timeout.

This frequency and timeout apply to all layer 7 service checks.

Example

The following command configures a flow redirect service-check frequency of 100 seconds and a timeout of 300 seconds:

```
config flow-redirect timer service-check frequency 100 timeout 300
```

History

This command was first available in ExtremeWare 6.2.1.

Platform Availability

This command is available on all "*i*" series platforms.

config flow-redirect timer tcp-port-check

```
config flow-redirect timer tcp-port-check frequency <seconds> timeout
<seconds>
```

Description

Configures the flow redirect TCP port check frequency and timeout.

Syntax Description

frequency	Specifies the tcp-port-check frequency. The range is 5 to 120.
timeout	Specifies the tcp-port-check timeout. The range is 5 to 300.

Default

The default frequency is 10 seconds.

The default timeout is 30 seconds.

Usage Guidelines

The frequency must be less than the timeout.

Example

The following command configures a flow redirect tcp-port-check frequency of 15 seconds and a timeout of 45 seconds:

```
config flow-redirect timer tcp-port-check frequency 15 timeout 45
```

History

This command was first available in ExtremeWare 6.2.1.

Platform Availability

This command is available on all "*i*" series platforms.

config slb esrp vlan

```
config slb esrp vlan <vlan name> [add | delete] unit [number]
```

Description

Configures all virtual servers with the specified unit number to match the state of the specified ESRP VLAN.

Syntax Description

vlan name	Specifies an ESRP VLAN.
unit	Specifies a unit identifier on a virtual server. The range is 1 to 16.

Default

The default is unit 1.

Usage Guidelines

You must configure ESRP for the VLAN that you specify.

Virtual servers added with a unit number that is already configured for ESRP failover automatically match the ESRP state configured for that unit number.

Use the unit number to associate a group of virual servers with an ESRP VLAN so that ESRP controls the failover state of the virtual servers. To set the unit number of a virtual server, use the following command:

```
config slb vip
```

For simplicity, Extreme Networks recommends that you put client, server, and virtual server VLANs in the same ESRP group.

Example

The following command configures ESRP VLAN “servers” to control the failover state of all virtual servers configured with unit 3:

```
config slb esrp vlan servers add unit 3
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb failover alive-frequency

```
config slb failover alive-frequency <seconds> timeout <seconds>
```

Description

Configures the frequency at which the local SLB device polls the remote SLB device.

Syntax Description

alive-frequency	The frequency at which the local SLB device polls the remote SLB device. The range is 1 to 60.
timeout	The amount of time within which the local switch must receive a response from the remote switch. The range is 1 to 60.

Default

The default alive frequency is 1 second.

The default timeout is 3 seconds.

Usage Guidelines

The frequency must be less than the timeout. Extreme Networks recommends that you set the timout greater than an even multiple of the frequency.

To enable active-active operation, use the following command:

```
enable slb failover
```

Example

The following command sets the alive frequency to 5 seconds and the timeout to 10 seconds:

```
config slb alive-frequency 5 timeout 10
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb failover dead-frequency

```
config slb failover dead-frequency <seconds>
```

Description

Configures the frequency at which the local switch attempts to re-establish communication with the unresponsive remote switch.

Syntax Description

dead-frequency	The frequency at which the local switch attempts to re-establish communication with the unresponsive remote switch. The range is 1 to 60.

Default

The default dead frequency is 2 seconds.

Usage Guidelines

To enable active-active operation, use the following command:

```
enable slb failover
```

Example

The following command sets the dead frequency to 5 seconds:

```
config slb dead-frequency 5
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb failover failback-now

```
config slb failover failback-now
```

Description

Configures the local SLB to release the remote SLB resources if the remote SLB is alive.

Syntax Description

This command has no arguments or variables.

Default

N/A

Usage Guidelines

When an active SLB unit fails and recovers, and manual failback is enabled, use this command to force the recovered SLB unit to become the active unit. Executing this command does not affect the SLB configuration.

To enable manual failback, use the following command:

```
enable slb failover manual-failback
```

To disable manual failback, use the following command:

```
disable slb failover manual-failback
```

Example

The following command forces SLB to immediately failback to the backup unit:

```
config slb failover failback-now
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb failover ping-check

```
config slb failover ping-check <ip address> {frequency <seconds> timeout
<seconds>}
```

Description

Configures the SLB device to actively determine if a remote gateway is reachable by performing a ping.

Syntax Description

ip address	Specifies the IP address of the remote gateway.
frequency	Specifies the frequency of pings sent to the remote gateway. The range is 1 to 60.
timeout	Specifies the time before the local device declares the remote gateway down. The range is 1 to 60.

Default

The default frequency is 1 second.

The default timeout is 3 seconds.

Usage Guidelines

The frequency must be less than the timeout.

If the external gateway is not reachable, the virtual servers failover to the remote SLB device.

Do not configure ping-check to the remote SLB switch. If you configure ping-check to the remote SLB switch and the remote switch fails, the local switch also fails.

Example

The following command sets the IP address of the remote gateway to 10.10.10.21 with a ping frequency of 5 seconds and a timeout of 10 seconds:

```
config slb failover ping-check 10.10.10.21 frequency 5 timeout 10
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb failover unit

```
config slb failover unit <number> remote-ipaddress <ip address>
local-ipaddress <ip address> {L4-port <port number>}
```

Description

Configures the switch for active-active operation.

Syntax Description

unit	Specifies a unit identifier on a virtual server. The range is 1 to 16.
number	Specifies a number from 1 - 16.
remote-ipaddress	Specifies the remote peer IP address.
local-ipaddress	Specifies the local failover IP address.
ip address	Specifies an IP address.
L4-port	Specifies the TCP port used for keep alive packets between failover peers.
port number	Specifies a port.

Default

The default L4-port is 1028.

Usage Guidelines

You must configure both active switches. You must use the actual IP address of the switches for the `remote-ip` and `local-ip`; you cannot use the IP address of a virtual server.

To enable active-active operation, use the following command:

```
enable slb failover
```

Extreme Networks recommends that you use a dedicated layer 2 VLAN to connect the two active-active switches.

Example

The following command configures the local SLB switch (with an IP address of 10.10.10.22) to direct unit 2 virtual servers to failover to the SLB switch with an IP address of 10.10.10.21:

```
config slb failover unit 2 remote-ipaddress 10.10.10.21 local-ipaddress 10.10.10.22
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb global connection-block

```
config slb global connection-block <number>
```

Description

Configures the number of SLB connections to allocate in memory, which improves performance.

Syntax Description

number	Specifies the number of connection blocks. The range is 100 to 20,000.

Default

The default is 10,000.

Usage Guidelines

Use this command when you are sure that you will have a minimum guaranteed number of connections. Additional connection blocks are allocated when necessary.

Do not use this command unless you are absolutely sure that you will use the memory allocated.

Example

The following command allocates memory for 500 connections:

```
config slb global connection-block 500
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb global connection-timeout

```
config slb global connection-timeout <seconds>
```

Description

Configures the connection timeout for transparent and translation modes.

Syntax Description

seconds	Specifies the number of seconds. The range is 1 to 180.

Default

The default is one second.

Usage Guidelines

None.

Example

The following command configures the connection timeout for 50 seconds:

```
config slb global connection-timeout 50
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb global ftp

```
config slb global ftp user <user name> {password {encrypted} <password>}
```

Description

Configures the default parameters for layer 7 FTP service checking.

Syntax Description

user	Specifies the user name that is checked.
password	Specifies the password for the user name.
encrypted	Encrypts the password stored in the configuration file.

Default

The default value for user and password is anonymous.

Usage Guidelines

If you do not enter a password, you are prompted for the password twice. Extreme Networks recommends that you use a password.

The FTP service check provides a more thorough check than ping check, because the FTP service check logs into the service.

To enable service checking, use the following command:

```
enable slb vip <vip name> service-check
```

To configure the frequency and timeout of service checks, use the following command:

```
config slb global service-check
```

To configure the parameters for a specific virtual server, use the following command:

```
config slb vip <vip name> service-check ftp
```

Example

The following command configures service check to login using the user name “service” and the password “check”:

```
config slb global ftp user service password check
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb global http

```
config slb global http url <url> match-string [<any-content> | alphanumeric
string]
```

Description

Configures the default parameters for layer 7 HTTP service checking.

Syntax Description

url	Specifies a URL.
match string	Specifies the text to be matched at the specified URL.
any-content	Specifies that any content confirms check.
alphanumeric string	Specifies the text string to match.

Default

The default value for url is /.

The default match string is any content.

Usage Guidelines

The HTTP service check provides a more thorough check than ping check, because the HTTP service check connects to a specific URL and checks for a specific text string. Extreme Networks recommends that you create a specific URL dedicated to this check.

Do not include “http://” in the URL. To check a URL beyond the root directory, include the path in the specified URL. The maximum length of a URL is 255 characters.

To enable service checking, use the following command:

```
enable slb vip <vip name> service-check
```

To configure the frequency and timeout of service checks, use the following command:

```
config slb global service-check
```

To configure the parameters for a specific virtual server, use the following command:

```
config slb vip <vip name> service-check ftp
```

Example

The following command configures service check to access http://www.checktest.com and look for the text “test”:

```
config slb global http url www.checktest.com match-string test
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb global nntp

```
config slb global nntp <newsgroup>
```

Description

Configures the default parameters for layer 7 NNTP service checking.

Syntax Description

newsgroup	Specifies a newsgroup.

Default

The default newsgroup is ebusiness.

Usage Guidelines

The NNTP service check provides a more thorough check than ping check, because the NNTP service check logs into the service.

To enable service checking, use the following command:

```
enable slb vip <vip name> service-check
```

To configure the frequency and timeout of service checks, use the following command:

```
config slb global service-check
```

To configure the parameters for a specific virtual server, use the following command:

```
config slb vip <vip name> service-check ftp
```

Example

The following command configures the service check to log into the newsgroup “comp.dcom.lans.ethernet”:

```
config slb global nntp comp.dcom.lans.ethernet
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb global persistence-level

```
config slb global persistence-level [same-vip-same-port | same-vip-any-port
| any-vip]
```

Description

Configures the persistence level globally.

Syntax Description

same-vip-same-port	Specifies that an entry must match both virtual server and port for persistence.
same-vip-any-port	Specifies that an entry must match virtual server, and can be any port.
any-vip	Specifies that an entry can match any port on any virtual server.

Default

The default level is `same-vip-same-port`.

Usage Guidelines

Use this command when different virtual servers do not require different persistence settings.

If you configure `any-vip` persistence, ensure that all virtual servers in all pools have the same services.

Example

The following command sets the global persistence level to `any-vip`:

```
config slb global persistence-level any-vip
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb global persistence-method

```
config slb global persistence-method [per-packet | per-session]
```

Description

Configures the behavior of the persistence timer.

Syntax Description

per-packet	Resets the persistence timer at the receipt of each packet.
per-session	Resets the persistence timer at the beginning of the session. When the timer expires, persistence for the session ends.

Default

The default method is `per-session`.

Usage Guidelines

Using per-packet persistence requires more CPU processing.

To set the persistence timer, use the following command:

`config slb vip <vip name> client-persistence-timeout`

Example

The following command sets the global persistence method to expire at the end of the session:

`config slb global persistence-method per-session`

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb global ping-check

```
config slb global ping-check frequency <seconds> timeout <seconds>
```

Description

Configures default health checking frequency and timeout period using layer 3-based pinging of the physical node.

Syntax Description

frequency	Specifies the frequency of the ping check. The range is 1 to 60 seconds.
timeout	Specifies the timeout of the ping check. The range is 1 to 60 seconds.

Default

The default frequency is 10 seconds.

The default timeout is 30 seconds.

Usage Guidelines

This command sets the global values for ping check. Use the global values if your servers are all equally reliable. You can configure a node to override the global values using the following command:

```
config slb node <ip address> ping-check
```

The frequency must be less than the timeout.

If the pinged node does not respond within the specified timeout period (three ping intervals by default), the node is considered down.

Shorter ping intervals require more CPU processing.

Example

The following command sets the global ping-check frequency to 5 seconds and the timeout to 15 seconds:

```
config slb global ping-check frequency 5 timeout 15
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb global pop3

```
config slb global pop3 user <user name> {password {encrypted} <password>}
```

Description

Configures the default parameters for layer 7 POP3 service checking.

Syntax Description

user	Specifies the user name that is checked.
password	Specifies the password for the user name.
encrypted	Encrypts the password stored in the configuration file.

Default

The default value for user and password is anonymous.

Usage Guidelines

If you do not enter a password, you are prompted for the password twice. Extreme Networks recommends that you use a password.

The POP3 service check provides a more thorough check than ping check, because the POP3 service check logs into the service.

To enable service checking, use the following command:

```
enable slb vip <vip name> service-check
```

To configure the frequency and timeout of service checks, use the following command:

```
config slb global service-check
```

To configure the parameters for a specific virtual server, use the following command:

```
config slb vip <vip name> service-check ftp
```

Example

The following command configures the service check to login using the user name “service” and the password “check”:

```
config slb global pop3 user service password check
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb global service-check

```
config slb global service-check frequency <seconds> timeout <seconds>
```

Description

Configures default health checking frequency and timeout period using layer 7-based application-dependent checking.

Syntax Description

frequency	Specifies the frequency of the service check. The range is 15 to 300 seconds.
timeout	Specifies the timeout of the service check. The range is 5 to 300 seconds.

Default

The default frequency is 60 seconds.

The default timeout is 180 seconds.

Usage Guidelines

The frequency must be less than the timeout.

If the health check frequency and timeout are not specified for a specific virtual server, the global values are used. To set specific frequency and timeout values for a virtual server, use the following command:

```
config slb vip <vip name> service-check
```

Shorter intervals require more CPU processing.

Example

The following command sets the service check frequency to 90 seconds and the timeout to 270 seconds:

```
config slb global service-check frequency 90 timeout 270
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb global smtp

```
config slb global smtp <dns domain>
```

Description

Configures the default parameters for layer 7 SMTP service checking.

Syntax Description

dns domain	Specifies the domain to check.

Default

The default value for `dns domain` is the switch's domain. If the switch does not have a DNS domain configured, the value is "mydomain.com".

Usage Guidelines

The SMTP service check provides a more thorough check than ping check, because the SMTP service check accesses the service.

To enable service checking, use the following command:

```
enable slb vip <vip name> service-check
```

To configure the frequency and timeout of service checks, use the following command:

```
config slb global service-check
```

To configure the parameters for a specific virtual server, use the following command:

```
config slb vip <vip name> service-check ftp
```

Example

The following command configures the service check to access the DNS domain servicecheck.domain.com:

```
config slb global smtp servicecheck.domain.com
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "i" series platforms.

config slb global synguard

```
config slb global synguard max-unacknowledged-SYNs <number>
```

Description

Configures the the SYN-guard feature.

Syntax Description

max-unacknowledged-SYNs	Specifies the number of half-open connections that the switch allows. The range is 10 to 4000.

Default

The default value is 50.

Usage Guidelines

If the number of half-open connections exceeds the number specified, the switch immediately ages out the half-open connections. This only applies to connections from the same source IP address.

SYN-guard is disabled by default. To enable SYN-guard, use the following command:

```
enable slb global synguard
```

SYN-guard is automatically enabled if you configure a `max-unacknowledged-SYNs` value greater than 0. A `max-unacknowledged-SYNs` value of 0 automatically disables SYN-guard.

Example

The following command configures the SYN-guard feature to age out half-open connections from the same source IP address when the number of connections exceeds 30:

```
config slb global synguard max-unacknowledged-SYNs 30
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb global tcp-port-check

```
config slb global tcp-port-check frequency <seconds> timeout <seconds>
```

Description

Configures default health checking frequency and timeout period using layer 4-based TCP port testing.

Syntax Description

frequency	Specifies the frequency of the TCP port check. The range is 5 to 120 seconds.
timeout	Specifies the timeout of the TCP port check. The range is 5 to 300 seconds.

Default

The default frequency is 30 seconds.

The default timeout is 90 seconds.

Usage Guidelines

The frequency must be less than the timeout.

The TCP port check is the least intrusive health check, as it does not log into or access the server.

If the frequency and timeout are not specified for a specific node, the global values are used. You can configure a node to override the global values using the following command:

```
config slb node <ip address> : <L4 port> tcp-port-check
```

To enable TCP port checking, use the following command:

```
enable slb node tcp-port-check
```

Example

The following command sets the global TCP-port-check frequency to 15 seconds and timeout to 45 seconds:

```
config slb global tcp-port-check frequency 15 timeout 45
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb global telnet

```
config slb global telnet userid <userid> password {encrypted} {<password>}
```

Description

Configures the default parameters for layer 7 telnet service checking.

Syntax Description

user	Specifies the user name that is checked.
password	Specifies the password for the user name.
encrypted	Encrypts the password stored in the configuration file.

Default

The default value for user and password is anonymous.

Usage Guidelines

If you do not enter a password, you are prompted for the password twice. Extreme Networks recommends that you use a password.

The telnet service check provides a more thorough check than ping check, because the telnet service check logs into the service.

To enable service checking, use the following command:

```
enable slb vip <vip name> service-check
```

To configure the frequency and timeout of service checks, use the following command:

```
config slb global service-check
```

To configure the parameters for a specific virtual server, use the following command:

```
config slb vip <vip name> service-check ftp
```

Example

The following command configures the service check to login using the user name “service” and the password “check”:

```
config slb global telnet user service password check
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb gogo-mode health-check

```
config slb gogo-mode <port number> health-check <ip address>
```

Description

Configures the health checker with the common IP addresses of the GoGo mode servers in this group.

Syntax Description

port number	Specifies the GoGo mode master port.
ip address	Specifies an IP address.

Default

N/A.

Usage Guidelines

Use this command to configure the IP address before configuring individual health checks.

Example

The following command configures the GoGo mode health check for the group with port 29 as the master port and an IP address of 192.168.200.2:

```
config slb gogo-mode 29 health-check 192.168.200.2
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all “i” series platforms.

config slb gogo-mode ping-check

```
config slb gogo-mode <port number> ping-check frequency <seconds> timeout
<seconds>
```

Description

Overrides the global default ping-check frequency and timeout values for this GoGo mode group.

Syntax Description

port number	Specifies the GoGo mode master port.
frequency	Specifies the frequency of the ping check. The range is 1 to 60 seconds.
timeout	Specifies the timeout of the ping check. The range is 1 to 60 seconds.

Default

The default frequency is 10 seconds.

The default timeout is 30 seconds.

Usage Guidelines

The frequency must be less than the timeout.

To restore a configured frequency and timeout back to the global default, specify 0 for the frequency and timeout.

Before you use this command, configure the IP address on the GoGo mode servers using the following command:

```
config slb gogo-mode health-check
```

To enable ping check for a GoGo mode group, use the following command:

```
enable slb gogo-mode <port number> ping-check
```

To disable ping check for a GoGo mode group, use the following command:

```
disable slb gogo-mode <port number> ping-check
```

Example

The following command configures a GoGo mode ping check frequency of 15 seconds and a timeout of 45 seconds for the group with port 29 as the master port:

```
config slb gogo-mode 29 ping-check frequency 15 timeout 45
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all “i” series platforms.

config slb gogo-mode service-check ftp

```
config slb gogo-mode <port number> service-check ftp {<L4-port>} {user
<user> | password {encrypted} <password>}
```

Description

Configures the FTP service check parameters for a GoGo mode group.

Syntax Description

port number	Specifies the GoGo mode master port.
L4-port	Specifies a layer 4 port.
user	Specifies the user name that is checked.
password	Specifies the password for the user name.
encrypted	Encrypts the password stored in the configuration file.

Default

The default value for user and password is anonymous.

Usage Guidelines

Before you use this command, configure the IP address on the GoGo mode servers using the following command:

```
config slb gogo-mode health-check
```

To enable service check for a GoGo mode group, use the following command:

```
enable slb gogo-mode <port number> service-check
```

To disable service check for a GoGo mode group, use the following command:

```
disable slb gogo-mode <port number> service-check
```

Example

The following command configures GoGo mode service check for the group with port 29 as the master port to login using the user name “service” and the password “check”:

```
config slb gogo-mode 29 service-check ftp user service password check
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all “i” series platforms.

config slb gogo-mode service-check http

```
config slb gogo-mode <port number> service-check http {<L4-port>} {url
<url> match-string [any-content | <alphanumeric string>]}
```

Description

Configures the HTTP service check parameters for a GoGo mode group.

Syntax Description

port number	Specifies the GoGo mode master port.
L4-port	Specifies a layer 4 port.
url	Specifies a URL.
any-content	Specifies that any content confirms check.
alphanumeric string	Specifies the text string to match.

Default

The default value for url is /.

The default match string is any content.

Usage Guidelines

This command accesses the specified URL and checks for the specified alphanumeric string in the first 1000 bytes. Extreme Networks recommends that you create a specific URL dedicated to this check.

Do not include “http://” in the URL. To check a URL beyond the root directory, include the path in the specified URL. The maximum length of a URL is 255 characters.

Before you use this command, configure the IP address on the GoGo mode servers using the following command:

`config slb gogo-mode health-check`

To enable service check for a GoGo mode group, use the following command:

`enable slb gogo-mode <port number> service-check`

To disable service check for a GoGo mode group, use the following command:

`disable slb gogo-mode <port number> service-check`

Example

The following command configures GoGo mode service check for the group with port 29 as the master port to access http://www.checktest.com and look for the text “test”:

```
config slb gogo-mode 29 service-check http url www.checktest.com match-string test
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all “i” series platforms.

config slb gogo-mode service-check pop3

```
config slb gogo-mode <port number> service-check pop3 {<L4-port>} {userid
<userid> | password {encrypted} <password>}
```

Description

Configures the service check parameters for a GoGo mode group.

Syntax Description

port number	Specifies the GoGo mode master port.
L4-port	Specifies a layer 4 port.
user	Specifies the user name that is checked.
password	Specifies the password for the user name.
encrypted	Encrypts the password stored in the configuration file.

Default

The default value for user and password is anonymous.

Usage Guidelines

Before you use this command, configure the IP address on the GoGo mode servers using the following command:

`config slb gogo-mode health-check`

To enable service check for a GoGo mode group, use the following command:

`enable slb gogo-mode <port number> service-check`

To disable service check for a GoGo mode group, use the following command:

`disable slb gogo-mode <port number> service-check`

Example

The following command configures GoGo mode service check for the group with port 29 as the master port to login using the user name “service” and the password “check”:

`config slb gogo-mode 29 service-check pop3 user service password check`

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all “i” series platforms.

config slb gogo-mode service-check smtp

```
config slb gogo-mode <port number> service-check smtp {<L4-port>} {<dns
domain>}
```

Description

Configures the service check parameters for a GoGo mode group.

Syntax Description

port number	Specifies the GoGo mode master port.
L4-port	Specifies a layer 4 port.
dns domain	Specifies the domain to check.

Default

The default value for `dns domain` is the switch's domain. If the switch does not have a DNS domain configured, the value is "mydomain.com".

Usage Guidelines

Before you use this command, configure the IP address on the GoGo mode servers using the following command:

`config slb gogo-mode health-check`

To enable service check for a GoGo mode group, use the following command:

`enable slb gogo-mode <port number> service-check`

To disable service check for a GoGo mode group, use the following command:

`disable slb gogo-mode <port number> service-check`

Example

The following command configures the GoGo mode service check for the group with port 29 as the master port to access the DNS domain servicecheck.domain.com:

`config slb gogo-mode 29 service-check smtp servicecheck.domain.com`

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all "i" series platforms.

config slb gogo-mode service-check telnet

```
config slb gogo-mode <port number> service-check telnet {<L4-port>} {user
<user name> | password {encrypted} <password>}
```

Description

Configures the service check parameters for a GoGo mode group.

Syntax Description

port number	Specifies the GoGo mode master port.
L4-port	Specifies a layer 4 port.
user	Specifies the user name that is checked.
password	Specifies the password for the user name.
encrypted	Encrypts the password stored in the configuration file.

Default

The default value for user and password is anonymous.

Usage Guidelines

Before you use this command, configure the IP address on the GoGo mode servers using the following command:

`config slb gogo-mode health-check`

To enable service check for a GoGo mode group, use the following command:

`enable slb gogo-mode <port number> service-check`

To disable service check for a GoGo mode group, use the following command:

`disable slb gogo-mode <port number> service-check`

Example

The following command configures GoGo mode service check for the group with port 29 as the master port to login using the user name “service” and the password “check”:

`config slb gogo-mode 29 service-check telnet user service password check`

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all “i” series platforms.

config slb gogo-mode service-check timer

```
config slb gogo-mode <port number> service-check timer [all | ftp | http |
telnet | smtp | nntp | pop3 | <TCP port number>] frequency <seconds>
timeout <seconds>
```

Description

Overrides the global service-check frequency and timeout values.

Syntax Description

port number	Specifies the GoGo mode master port.
all	Specifies all service checks.
ftp	Specifies the FTP service check.
http	Specifies the HTTP service check.
telnet	Specifies the telnet service check.
smtp	Specifies the SMTP service check.
nntp	Specifies the NNTP service check.
pop3	Specifies the POP3 service check.
TCP port number	Specifies a TCP port, instead of a service, for the service check.
frequency	Specifies the frequency of the service check. The range is 15 to 300 seconds.
timeout	Specifies the timeout of the service check. The range is 15 to 300 seconds.

Default

N/A.

Usage Guidelines

You can use this command at any time. This command affects the frequency and timeout for the specified service-check in the specified GoGo mode group.

The frequency must be less than the timeout.

Example

The following command configures GoGo mode FTP service check for the group with port 29 as the master port with a frequency of 15 seconds and a timeout of 45 seconds:

```
config slb gogo-mode 29 service-check timer ftp frequency 15 timeout 45
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all “i” series platforms.

config slb gogo-mode tcp-port-check add

```
config slb gogo-mode <port number> tcp-port-check add [ftp | http | https |
imap4 | ldap | nntp | pop3 | smtp | socks | telnet | tftp | web | www |
<TCP port number>]
```

Description

Adds the specified layer 4 port.

Syntax Description

port number	Specifies the GoGo mode master port.
ftp	Specifies the FTP TCP-port-check.
http	Specifies the HTTP TCP-port-check.
https	Specifies the HTTPS TCP-port-check.
imap4	Specifies the IMAP4 TCP-port-check.
ldap	Specifies the LDAP TCP-port-check.
nntp	Specifies the NNTP TCP-port-check.
pop3	Specifies the POP3 TCP-port-check.
smtp	Specifies the SMTP TCP-port-check.
socks	Specifies the SOCKS TCP-port-check.
telnet	Specifies the telnet TCP-port-check.
tftp	Specifies the TFTP TCP-port-check.
web	Specifies the Web TCP-port-check.
www	Specifies the www TCP-port-check.
TCP port number	Specifies a TCP port for the TCP-port-check.

Default

N/A.

Usage Guidelines

This command adds the port to the specified TCP-port-check in the specified GoGo mode group. You can only add a single port with each command; to add multiple ports, you must enter multiple commands.

Example

The following command adds FTP as a GoGo mode TCP-port-check for the group with port 29 as the master port:

```
config slb gogo-mode 29 tcp-port-check add ftp
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all “i” series platforms.

config slb gogo-mode tcp-port-check delete

```
config slb gogo-mode <port number> tcp-port-check delete [ftp | http |
https | imap4 | ldap | nntp | pop3 | smtp | socks | telnet | tftp | web |
www | <TCP port number>]
```

Description

Deletes the specified layer 4 port.

Syntax Description

port number	Specifies the GoGo mode master port.
ftp	Specifies the FTP TCP-port-check.
http	Specifies the HTTP TCP-port-check.
https	Specifies the HTTPS TCP-port-check.
imap4	Specifies the IMAP4 TCP-port-check.
ldap	Specifies the LDAP TCP-port-check.
nntp	Specifies the NNTP TCP-port-check.
pop3	Specifies the POP3 TCP-port-check.
smtp	Specifies the SMTP TCP-port-check.
socks	Specifies the SOCKS TCP-port-check.
telnet	Specifies the telnet TCP-port-check.
tftp	Specifies the TFTP TCP-port-check.
web	Specifies the Web TCP-port-check.
www	Specifies the www TCP-port-check.
TCP port number	Specifies a TCP port for the TCP-port-check.

Default

N/A.

Usage Guidelines

This command deletes the port from the specified TCP-port-check in the specified GoGo mode group. You can only delete a single port with each command; to delete multiple ports, you must enter multiple commands.

Example

The following command deletes FTP from the GoGo mode TCP-port-check for the group with port 29 as the master port:

```
config slb gogo-mode 29 tcp-port-check delete ftp
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all “i” series platforms.

config slb gogo-mode tcp-port-check timer

```
config slb gogo-mode <port number> tcp-port-check timer [ftp | http | https
| imap4 | ldap | nntp | pop3 | smtp | socks | telnet | tftp | web | www |
<TCP port number>] frequency <seconds> timeout <seconds>
```

Description

Overrides the global TCP-port-check frequency and timeout values.

Syntax Description

port number	Specifies the GoGo mode master port.
ftp	Specifies the FTP TCP-port-check.
http	Specifies the HTTP TCP-port-check.
https	Specifies the HTTPS TCP-port-check.
imap4	Specifies the IMAP4 TCP-port-check.
ldap	Specifies the LDAP TCP-port-check.
nntp	Specifies the NNTP TCP-port-check.
pop3	Specifies the POP3 TCP-port-check.
smtp	Specifies the SMTP TCP-port-check.
socks	Specifies the SOCKS TCP-port-check.
telnet	Specifies the telnet TCP-port-check.
tftp	Specifies the TFTP TCP-port-check.
web	Specifies the Web TCP-port-check.
www	Specifies the www TCP-port-check.
TCP port number	Specifies a TCP port for the TCP-port-check.
frequency	Specifies the frequency of the TCP port check. The range is 5 to 120 seconds.
timeout	Specifies the timeout of the TCP port check. The range is 5 to 300 seconds.

Default

N/A.

Usage Guidelines

This command affects only the specified GoGo mode group.

To set the global TCP-port-check frequency and timeout, use the following command:

`config slb global tcp-port-check`

The frequency must be less than the timeout.

Example

The following command configures GoGo mode FTP TCP-port-check for the group with port 29 as the master port with a frequency of 15 seconds and a timeout of 45 seconds:

```
config slb gogo-mode 29 tcp-port-check timer ftp frequency 15 timeout 45
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all “i” series platforms.

config slb L4-port

```
config slb L4-port [ftp | http | https | imap4 | ldap | nntp | pop3 | smtp
| socks | telnet | tftp | web | www | <TCP or UDP port number>]
[treaper-timeout <seconds> } udp-idle-timeout <seconds>]
```

Description

Configures the inactive period for TCP or UDP before the connection is aged out.

Syntax Description

ftp	Specifies the FTP service.
http	Specifies the HTTP service.
https	Specifies the HTTPS service.
imap4	Specifies the IMAP4 service.
ldap	Specifies the LDAP service.
nntp	Specifies the NNTP service.
pop3	Specifies the POP3 service.
smtp	Specifies the SMTP service.
socks	Specifies the SOCKS service.
telnet	Specifies the telnet service.
tftp	Specifies the TFTP service.
web	Specifies the Web service.
www	Specifies the www service
TCP or UDP port number	Specifies a TCP or UDP port for the service.
treaper-timeout	Specifies the timeout for TCP. The range is to .
udp-idle-timeout	Specifies the timeout for UDP. The range is to .

Default

The default `treaper-timeout` is .

The default `udp-idle-timeout` is .

Usage Guidelines

You must configure the port and add it to a pool before you use this command. The timeout value affects all connections to the specified service on all virtual servers.

To set the timeout values for a wildcard virtual server, use a TCP or UDP port number of 0.

Example

The following command configures the ftp nodes with a TCP idle period of 30 seconds:

```
config slb l4-port ftp treaper-timeout 30
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb node max-connections

```
config slb node <ip address>:[ftp | http | https | imap4 | ldap | nntp |
pop3 | smtp | socks | telnet | tftp | web | www | <TCP or UDP port number>]
max-connections <number>
```

Description

Configures the maximum number of simultaneous connections that can be established to a node.

Syntax Description

ip address	Specifies the IP address of the node.
ftp	Specifies the FTP TCP-port-check.
http	Specifies the HTTP TCP-port-check.
https	Specifies the HTTPS TCP-port-check.
imap4	Specifies the IMAP4 TCP-port-check.
ldap	Specifies the LDAP TCP-port-check.
nntp	Specifies the NNTP TCP-port-check.
pop3	Specifies the POP3 TCP-port-check.
smtp	Specifies the SMTP TCP-port-check.
socks	Specifies the SOCKS TCP-port-check.
telnet	Specifies the telnet TCP-port-check.
tftp	Specifies the TFTP TCP-port-check.
web	Specifies the Web TCP-port-check.
www	Specifies the www TCP-port-check.
TCP or UDP port number	Specifies a TCP or UDP port for the TCP-port-check.
max-connections	Specifies the maximum number of simultaneous connections. The range is 0 to 999999999.

Default

The default is 0.

Usage Guidelines

Use this command to limit the number of connections possible to a server with limited capabilities. Use `max-connections` of 0 to specify no limit.

Example

The following command configures the server with an IP address of 10.1.1.2:80 to accept a maximum of 10 connections:

```
config slb node 10.1.1.2 : 80 max-connections 10
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb node ping-check

```
config slb node <ip address> ping-check frequency <seconds> timeout
<seconds>
```

Description

Overrides the global default frequency and timeout values for this node.

Syntax Description

ip address	Specifies the IP address of the node.
frequency	Specifies the frequency of the ping check. The range is 1 to 60 seconds.
timeout	Specifies the timeout of the ping check. The range is 1 to 60 seconds.

Default

N/A.

Usage Guidelines

The frequency must be less than the timeout.

To set the global ping-check frequency and timeout, use the following command:

```
config slb global ping-check
```

Example

The following command sets the ping-check for the node with an IP address of 10.2.1.2 to a frequency of 30 seconds and a timeout of 90 seconds:

```
config slb node 10.2.1.2 ping-check frequency 30 timeout 90
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb node tcp-port-check

```
config slb node <ip address>:[ftp | http | https | imap4 | ldap | nntp |
pop3 | smtp | socks | telnet | tftp | web | www | <TCP or UDP port number>]
tcp-port-check frequency <seconds> timeout <seconds>
```

Description

Overrides the global default frequency and timeout values for this node.

Syntax Description

ip address	Specifies the IP address of the node.
ftp	Specifies the FTP TCP-port-check.
http	Specifies the HTTP TCP-port-check.
https	Specifies the HTTPS TCP-port-check.
imap4	Specifies the IMAP4 TCP-port-check.
ldap	Specifies the LDAP TCP-port-check.
nntp	Specifies the NNTP TCP-port-check.
pop3	Specifies the POP3 TCP-port-check.
smtp	Specifies the SMTP TCP-port-check.
socks	Specifies the SOCKS TCP-port-check.
telnet	Specifies the telnet TCP-port-check.
tftp	Specifies the TFTP TCP-port-check.
web	Specifies the Web TCP-port-check.
www	Specifies the www TCP-port-check.
TCP or UDP port number	Specifies a TCP or UDP port for the TCP-port-check.
frequency	Specifies the frequency of the TCP port check. The range is 5 to 120 seconds.
timeout	Specifies the timeout of the TCP port check. The range is 5 to 300 seconds.

Default

N/A.

Usage Guidelines

To set the global TCP-port-check frequency and timeout, use the following command:

```
config slb global tcp-port-check
```

The frequency must be less than the timeout.

Example

The following command sets the FTP TCP-port-check for the node with an IP address of 10.2.1.2 to a frequency of 30 seconds and a timeout of 90 seconds:

```
config slb node 10.2.1.2 : ftp tcp-port-check frequency 30 timeout 90
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb pool add

```
config slb pool <pool name> add <ip address>:[ftp | http | https | imap4 |
ldap | nntp | pop3 | smtp | socks | telnet | tftp | web | www | <TCP or UDP
port number>] {ratio <number> | priority <number>}
```

Description

Adds a node to a pool.

Syntax Description

pool name	Specifies a pool.
ip address	Specifies the IP address of the node.
ftp	Specifies an FTP node.
http	Specifies an HTTP node.
https	Specifies an HTTPS node.
imap4	Specifies an IMAP4 node.
ldap	Specifies an LDAP node.
nntp	Specifies an NNTP node.
pop3	Specifies a POP3 node.
smtp	Specifies an SMTP node.
socks	Specifies a SOCKS node.
telnet	Specifies a telnet node.
tftp	Specifies a TFTP node.
web	Specifies a Web node.
www	Specifies a www node
TCP or UDP port number	Specifies a TCP or UDP port for the node.
ratio	Specifies the ratio for the ratio load balancing method. The range is 0 to 65,535.
priority	Specifies the priority for the priority load balancing method. The range is 1 to 65535.

Default

The default ratio is 1.

Usage Guidelines

This command also configures the ratio or priority for the ratio and priority load balancing methods.

You must create the pool before you add nodes. When you add a new node, ping-check is automatically enabled.

A ratio of 2 results in twice as much traffic as a ratio of 1. If all nodes use the same ratio, connections are distributed equally among the nodes. A ratio of 0 results in no traffic to the node. When you configure the ratio, use the smallest common denominator. For example, to configure a ratio of 25% and 75%, use ratios of 1 and 3, instead of 25 and 75.

To configure a pool to use the ratio load balancing method, use the following command:

```
config slb pool <pool name> lb-method ratio
```

Higher priority numbers indicate higher priority. To configure a pool to use the priority load balancing method, use the following command:

```
config slb pool <pool name> lb-method priority
```

To change the ratio or priority of a node that is already in a pool, use the following command:

```
config slb pool <pool name> member
```

Example

The following command adds the FTP node with an IP address of 10.2.1.2 to the pool “ftp” and configures the node with a priority of 2:

```
config slb pool ftp add 10.2.1.2 : ftp priority 2
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb pool delete

```
config slb pool <pool name> delete <ip address>:[ftp | http | https | imap4
| ldap | nntp | pop3 | smtp | socks | telnet | tftp | web | www | <TCP or
UDP port number>]
```

Description

Deletes a node from a pool.

Syntax Description

pool name	Specifies a pool.
ip address	Specifies the IP address of the node.
ftp	Specifies an FTP node.
http	Specifies an HTTP node.
https	Specifies an HTTPS node.
imap4	Specifies an IMAP4 node.
ldap	Specifies an LDAP node.
nntp	Specifies an NNTP node.
pop3	Specifies a POP3 node.
smtp	Specifies an SMTP node.
socks	Specifies a SOCKS node.
telnet	Specifies a telnet node.
tftp	Specifies a TFTP node.
web	Specifies a Web node.
www	Specifies a www node
TCP or UDP port number	Specifies a TCP or UDP port for the node.

Default

N / A.

Usage Guidelines

Deleting a node from a pool does not delete the node from other pools. You can delete all nodes in a pool by deleting the pool. To delete a pool, use the following command:

```
delete slb pool
```

Example

The following command deletes the FTP node with an IP address of 10.2.1.2 from the pool “ftp”:

```
config slb pool ftp delete 10.2.1.2 : ftp
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb pool lb-method

```
config slb pool <pool name> lb-method [round-robin | ratio | priority |
least-connections]
```

Description

Configures the SLB load balancing method.

Syntax Description

pool name	Specifies a pool.
round-robin	Specifies the round robin load balancing method.
ratio	Specifies the ratio load balancing method.
priority	Specifies the priority load balancing method.
least-connections	Specifies the least connections load balancing method.

Default

N / A.

Usage Guidelines

Use this command to change the load balancing method after you have already created a pool.

To set the ratio or priority of a node, use the following command:

```
config slb pool <pool name> member
```

Example

The following command changes the load balancing method for the pool “ftp” to ratio:

```
config slb pool ftp lb-method ratio
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb pool member

```
config slb pool <pool name> member <ip address>:[ftp | http | https | imap4
| ldap | nntp | pop3 | smtp | socks | telnet | tftp | web | www | <TCP or
UDP port number>] [ratio <number> | priority <number>]
```

Description

Configures the ratio or priority of an existing pool member.

Syntax Description

pool name	Specifies a pool.
ip address	Specifies the IP address of the node.
ftp	Specifies an FTP node.
http	Specifies an HTTP node.
https	Specifies an HTTPS node.
imap4	Specifies an IMAP4 node.
ldap	Specifies an LDAP node.
nntp	Specifies an NNTP node.
pop3	Specifies a POP3 node.
smtp	Specifies an SMTP node.
socks	Specifies a SOCKS node.
telnet	Specifies a telnet node.
tftp	Specifies a TFTP node.
web	Specifies a Web node.
www	Specifies a www node
TCP or UDP port number	Specifies a TCP or UDP port for the node.
ratio	Specifies the ratio for the ratio load balancing method. The range is 0 to 65,535.
priority	Specifies the priority for the priority load balancing method. The range is 1 to 65535.

Default

N / A.

Usage Guidelines

Use this command to change the ratio or priority of an existing node. To add a node to a pool (and set the ratio or priority), use the following command:

```
config slb pool <pool name> add
```

Example

The following command changes the priority of the FTP node with an IP address of 10.2.1.2 in the pool “ftp” to 2:

```
config slb pool ftp member 10.2.1.2 : ftp priority 2
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb proxy-client-persistence

```
config slb proxy-client-persistence [add | delete] <ip address>/<netmask>
```

Description

Configures a client subnet that should be treated as one persistent entity.

Syntax Description

ip address/netmask	Specifies an IP address and netmask.

Default

N/A.

Usage Guidelines

Use this command to force all clients from the specified proxy array to connect to the same physical server.

Example

The following command specifies that the subnet 10.10.10.20/24 should be treated as a single, persistent entity:

```
config slb proxy-client-persistence add 10.10.10.20/24
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb vip

```
config slb vip <vip name> unit [number]
```

Description

Configures the unit number for active-active failover.

Syntax Description

vip name	Specifies a virtual server.
unit	Specifies a unit identifier on a virtual server. The range is 1 to 16.

Default

The default unit is 1.

Usage Guidelines

You must first create the virtual server before you use this command. To create a virtual server, use the following command:

```
creat slb vip
```

Example

The following command configures the virtual server “test” with a unit number of 3:

```
config slb vip test unit 3
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb vip client-persistence-timeout

```
config slb vip [<vip name> | all] client-persistence-timeout <seconds>
```

Description

Configures the client persistence timeout value.

Syntax Description

vip name	Specifies a virtual server.
all	Specifies all virtual servers.
client-persistence-timeout	Specifies the persistence timeout. The range is 1 to 999,999,999.

Default

The default `client-persistence-timeout` is 3600.

Usage Guidelines

Extreme Networks recommends that you specify a short client persistence timeout, because longer timeout values consume more memory.

Example

The following command configures the virtual server “ftp” with a client persistence timeout of 3000 seconds:

```
config slb vip ftp client-persistence-timeout 3000
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb vip max-connections

```
config slb vip <vip name> max-connections <number>
```

Description

Configures the maximum connections allowed to a particular virtual server.

Syntax Description

vip name	Specifies a virtual server.
max-connections	Specifies the maximum number of connections allowed to a virtual server. The range is 0 to 999,999,999.

Default

The default value is 0.

Usage Guidelines

A value of 0 indicates that no maximum is enforced. When the maximum number of connections is reached, the server stops responding to new requests; existing connections are maintained.

Example

The following command sets the maximum connections to the virtual server “ftp” to 10:

```
config slb vip ftp max-connections 10
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb vip service-check frequency

```
config slb vip <vip name> service-check frequency <seconds> timeout
<seconds>
```

Description

Configures the layer 7 service check frequency and timeout for a particular virtual server.

Syntax Description

vip name	Specifies a virtual server.
frequency	Specifies the frequency of the service check. The range is 15 to 300 seconds.
timeout	Specifies the timeout of the service check. The range is 5 to 300 seconds.

Default

N/A.

Usage Guidelines

The frequency must be less than the timeout.

To return to the global values, specify 0 for frequency and timeout. To set the global service check frequency and timeout, use the following command:

```
config slb global service-check
```

Example

The following command sets the service check frequency to 15 and timeout to 45 for the virtual server "ftp":

```
config slb vip ftp service-check frequency 15 timeout 45
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "i" series platforms.

config slb vip service-check ftp

```
config slb vip <vip name> service-check ftp {user <user name> password
{encrypted} <password>}
```

Description

Configures layer 7 FTP service checking for a specific virtual server.

Syntax Description

vip name	Specifies a virtual server.
user	Specifies the user name that is checked.
password	Specifies the password for the user name.
encrypted	Encrypts the password stored in the configuration file.

Default

N / A.

Usage Guidelines

This command automatically enables service checking.

If you do not enter a password, you are prompted for the password twice. Extreme Networks recommends that you use a password.

The FTP service check provides a more thorough check than ping check, because the FTP service check logs into the service.

To configure the frequency and timeout of service checks, use the following command:

```
config slb global service-check
```

To configure the global parameters, use the following command:

```
config slb global ftp
```

Example

The following command configures service check to login using the user name “service” and the password “check” on the virtual server “ftpvip”:

```
config slb vip ftpvip service-check ftp user service password check
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb vip service-check http

```
config slb vip <vip name> service-check http {url <url> match-string
[any-content | <alphanumeric string>]}
```

Description

Configures layer 7 HTTP service checking for a specific virtual server.

Syntax Description

vip name	Specifies a virtual server.
url	Specifies a URL.
match string	Specifies the text to be matched at the specified URL.
any-content	Specifies that any content confirms check.
alphanumeric string	Specifies the text string to match.

Default

N/A.

Usage Guidelines

This command automatically enables service checking.

The HTTP service check provides a more thorough check than ping check, because the HTTP service check connects to a specific URL and checks for a specific text string. Extreme Networks recommends that you create a specific URL dedicated to this check.

Do not include "http://" in the URL. To check a URL beyond the root directory, include the path in the specified URL. The maximum length of a URL is 255 characters.

To configure the frequency and timeout of service checks, use the following command:

```
config slb global service-check
```

To configure the global parameters, use the following command:

```
config slb global http
```

Example

The following command configures service check to access http://www.checktest.com and look for the text "test" on the virtual server "httpvip":

```
config slb vip httpvip service-check http url www.checktest.com match-string test
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb vip service-check nntp

```
config slb vip <vip name> service-check nntp <newsgroup>
```

Description

Configures layer 7 NNTP service checking for a specific virtual server.

Syntax Description

vip name	Specifies a virtual server.
newsgroup	Specifies a newsgroup.

Default

N/A.

Usage Guidelines

This command automatically enables service checking.

The NNTP service check provides a more thorough check than ping check, because the NNTP service check logs into the service.

To configure the frequency and timeout of service checks, use the following command:

```
config slb global service-check
```

To configure the global parameters, use the following command:

```
config slb global nntp
```

Example

The following command configures the service check to log into the newsgroup "comp.dcom.lans.ethernet" on the virtual server "nntpvip":

```
config slb vip nntpvip service-check nntp comp.dcom.lans.ethernet
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "i" series platforms.

config slb vip service-check pop3

```
config slb vip <vip name> service-check pop3 user <user name> password
{encrypted} {password}
```

Description

Configures layer 7 POP3 service checking for a specific virtual server.

Syntax Description

vip name	Specifies a virtual server.
user	Specifies the user name that is checked.
password	Specifies the password for the user name.
encrypted	Encrypts the password stored in the configuration file.

Default

N / A.

Usage Guidelines

This command automatically enables service checking.

If you do not enter a password, you are prompted for the password twice. Extreme Networks recommends that you use a password.

The POP3 service check provides a more thorough check than ping check, because the POP3 service check logs into the service.

To configure the frequency and timeout of service checks, use the following command:

```
config slb global service-check
```

To configure the global parameters, use the following command:

```
config slb global pop3
```

Example

The following command configures the service check to login using the user name “service” and the password “check” to the virtual server “pop3vip”:

```
config slb vip pop3vip service-check pop3 user service password check
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb vip service-check smtp

```
config slb vip <vip name> service-check smtp {<dns domain>}
```

Description

Configures layer 7 SMTP service checking for a specific virtual server.

Syntax Description

vip name	Specifies a virtual server.
dns domain	Specifies the domain to check.

Default

N / A.

Usage Guidelines

This command automatically enables service checking.

The SMTP service check provides a more thorough check than ping check, because the SMTP service check accesses the service.

To configure the frequency and timeout of service checks, use the following command:

```
config slb global service-check
```

To configure the global parameters, use the following command:

```
config slb global smtp
```

Example

The following command configures the service check to access the DNS domain servicecheck.domain.com on the virtual server “smtpvip”:

```
config slb vip smtpvip service-check smtp servicecheck.domain.com
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config slb vip service-check telnet

```
config slb vip <vip name> service-check telnet {user <user name> password
{encrypted} <password>}
```

Description

Configures layer 7 telnet service checking for a specific virtual server.

Syntax Description

vip name	Specifies a virtual server.
user	Specifies the user name that is checked.
password	Specifies the password for the user name.
encrypted	Encrypts the password stored in the configuration file.

Default

N / A.

Usage Guidelines

This command automatically enables service checking.

If you do not enter a password, you are prompted for the password twice. Extreme Networks recommends that you use a password.

The telnet service check provides a more thorough check than ping check, because the telnet service check logs into the service.

To configure the frequency and timeout of service checks, use the following command:

`config slb global service-check`

To configure the global parameters, use the following command:

`config slb global telnet`

Example

The following command configures the service check to login using the user name “service” and the password “check” on the virtual server “telnetvip”:

`config slb vip telnetvip service-check telnet user service password check`

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

config vlan slb-type

```
config vlan <vlan name> slb-type [server | client | both | none]
```

Description

Marks a VLAN as either a server VLAN or a client VLAN.

Syntax Description

server	Configures the VLAN as a server VLAN.
client	Configures the VLAN as a client VLAN.
both	Configures the VLAN as both a server and a client VLAN.
none	Disables SLB on the VLAN.

Default

The default is none.

Usage Guidelines

Use the `both` option if a server originates or could possibly originate connections to other servers.

Example

The following command configures the VLAN “client_vlan” as a client VLAN:

```
config vlan client_vlan slb-type client
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

create flow-redirect

```
create flow-redirect <flow redirect> [any | tcp | tup | udp] destination
[<ip address> / <mask> | any] ip-port [<port> | any] source [<ip address> /
<mask> | any]
```

Description

Creates a flow redirect policy.

Syntax Description

flow redirect	Specifies a flow redirect policy.
any	Forwards traffic using either TCP or UDP on any IP port.
tcp	Forwards TCP traffic on a single IP port.
tup	Forwards traffic using either TCP or UDP on a single IP port.
udp	Forwards traffic using only UDP on a single IP port.
ip address	Specifies an IP address.
port	Specifies the port.

Default

N / A.

Usage Guidelines

Creating a flow redirect policy automatically enables flow redirect.

To delete a flow redirect policy, use the following command:

```
delete flow-redirect <flow redirect>
```

To rename or modify a flow redirect policy, you must delete and recreate the flow redirect policy.

Example

The following command creates a flow redirect policy named “http” that forwards TCP traffic to 10.1.1.10 port 80 from any source IP address:

```
create flow-redirect http tcp destination 10.1.1.10/29 ip-port 80 source any
```

History

This command was available in ExtremeWare 6.1.4. This command was modified in 6.2 to add the `tup` parameter.

Platform Availability

This command is available on all “*i*” series platforms.

create slb pool

```
create slb pool <pool name> {lb-method [least-connections | priority |
ratio | round-robin]}
```

Description

Creates a server pool and optionally assigns a load-balancing method to the pool.

Syntax Description

pool name	Specifies a pool.
lb-method	Specifies the load-balancing method.

Default

The default load-balancing method is round-robin.

Usage Guidelines

To change the load-balancing method of an existing pool, use the following command:

```
config slb pool <pool name> lb-method
```

To add a node to the pool (and set the ratio or priority), use the following command:

```
config slb pool <pool name> add
```

Example

The following command creates the pool “ftp_pool” and assigns the priority load-balancing method:

```
config slb pool ftp_pool lb-method priority
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

create slb vip

```
create slb vip <vip name> pool <pool name> mode [transparent | translation
| port-translation] <ip address> {- <upper range>} : <L4 port> {unit
<number>}
```

Description

Creates one or more new virtual servers.

Syntax Description

vip name	Specifies a virtual server.
pool name	Specifies a pool.
mode	Specifies the forwarding mode.
ip address	Specifies the IP address of the virtual server.
upper range	Specifies the upper IP address for a range of IP addresses.
L4 port	Specifies a port.
unit	Specifies a unit identifier on a virtual server. The range is 1 to 16.

Default

N/A.

Usage Guidelines

You must create the pool before assigning a virtual server to the pool. To create a pool, use the following command:

`create slb pool`

Example

The following command creates the virtual server “ftp_vip” with an IP address of 10.10.10.2 in the pool “ftp_pool” and assigns the port-translation forwarding mode:

```
config slb vip ftp_vip pool ftp_pool mode port-translation 10.10.10.2 : ftp
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

delete flow-redirect

```
delete flow-redirect <flow redirect>
```

Description

Deletes a flow redirect policy.

Syntax Description

flow redirect	Specifies a flow redirect policy.

Default

N/A.

Usage Guidelines

To rename or modify a flow redirect policy, you must delete and recreate the flow redirect policy.

Example

The following command deletes a flow redirect policy named “http”:

```
delete flow-redirect http
```

History

This command was available in ExtremeWare 6.1.4.

Platform Availability

This command is available on all “*i*” series platforms.

delete slb pool

```
delete slb pool [<pool name> | all]
```

Description

Deletes a server pool.

Syntax Description

pool name	Specifies a pool.
all	Specifies all pools.

Default

N / A.

Usage Guidelines

You must first delete all virtual servers before deleting the pool. To delete a virtual server, use the following command:

`delete slb vip`

Example

The following command the pool named “http_pool”:

`delete slb pool http_pool`

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

delete slb vip

```
delete slb vip [<vip name> | all]
```

Description

Deletes one or all virtual servers.

Syntax Description

vip name	Specifies a virtual server.
all	Specifies all virtual servers.

Default

N / A.

Usage Guidelines

You must use this command to delete all virtual servers from a pool before deleting the pool.

Example

The following command the virtual server named “http_vip”:

```
delete slb pool http_vip
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

disable flow-redirect

```
disable flow-redirect [all | <flow redirect>]
```

Description

Disables flow redirect.

Syntax Description

all	Specifies all flow policies.
flow redirect	Specifies a single flow redirect policy.

Default

The default parameter is all.

Flow redirect is disabled by default.

Usage Guidelines

When you create a new flow redirect policy, flow redirect is automatically enabled.

To enable flow redirect, use the following command:

```
enable flow-redirect
```

Example

The following command disables flow redirect for all flow policies:

```
disable flow-redirect all
```

History

This command was available in ExtremeWare 6.1.4.

Platform Availability

This command is available on all "*i*" series platforms.

disable slb

```
disable slb
```

Description

Disables SLB processing.

Syntax Description

This command has no arguments or variables.

Default

SLB is disabled by default.

Usage Guidelines

Disabling SLB causes the following to occur:

- Closes all connections.
- Withdraws virtual server routes or routes that do not respond with proxy ARP responses of virtual server addresses.
- Disconnects the switch from redundant SLB switches.

To enable SLB, use the following command:

```
enable slb
```

Example

The following command disables SLB:

```
disable slb
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

disable slb 3dns

```
disable slb 3dns iquery-client
```

Description

Disables 3DNS support.

Syntax Description

This command has no arguments or variables.

Default

3DNS is disabled by default.

Usage Guidelines

To enable 3DNS, use the following command:

```
enable slb 3dns iquery-client
```

Example

The following command disables 3DNS:

```
disable slb 3dns iquery-client
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

disable slb failover

```
disable slb failover
```

Description

Disables the SLB failover mechanism.

Syntax Description

This command has no arguments or variables.

Default

SLB failover is disabled by default.

Usage Guidelines

To enable SLB failover, use the following command:

`enable slb failover`

Example

The following command disables SLB failover:

`disable slb failover`

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

disable slb failover manual-failback

```
disable slb failover manual-failback
```

Description

Disables manual failback.

Syntax Description

This command has no arguments or variables.

Default

Manual failback is disabled by default.

Usage Guidelines

To enable manual failback, use the following command:

`enable slb failover manual-failback`

Example

The following command disables manual failback:

`disable slb failover manual-failback`

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

disable slb failover ping-check

```
disable slb failover ping-check
```

Description

Disables ping-check to an external gateway.

Syntax Description

This command has no arguments or variables.

Default

Ping-check is disabled by default.

Usage Guidelines

To enable ping-check, use the following command:
`enable slb failover ping-check`

Example

The following command disables ping-check:
`disable slb failover ping-check`

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

disable slb global synguard

```
disable slb global synguard
```

Description

Disables the TCP SYN-guard feature.

Syntax Description

This command has no arguments or variables.

Default

SYN-guard is disabled by default.

Usage Guidelines

To enable SYN-guard, use the following command:

```
enable slb global synguard
```

Example

The following command disables SYN-guard:

```
disable slb global synguard
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

disable slb gogo-mode

```
disable slb gogo-mode <port number>
```

Description

Disables GoGo mode processing.

Syntax Description

port number	Specifies the GoGo mode master port.

Default

GoGo mode is disabled by default.

Usage Guidelines

Before you disable GoGo mode, disconnect the servers, as they all have identical MAC and IP addresses, which can cause VLAN conflicts.

To enable GoGo mode, use the following command:

```
enable slb gogo-mode
```

Example

The following command disables GoGo mode for the group with port 29 as the master port:

```
disable slb gogo-mode 29
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

disable slb gogo-mode ping-check

```
disable slb gogo-mode <port number> ping-check
```

Description

Disables layer-3 ping-check to this GoGo mode group.

Syntax Description

port number	Specifies the GoGo mode master port.

Default

GoGo mode ping check is disabled by default.

Usage Guidelines

To enable ping-check for a GoGo mode group, use the following command:

`enable slb gogo-mode <port number> ping-check`

Example

The following command disables GoGo mode ping-check for the group with port 29 as the master port:

`disable slb gogo-mode 29 ping-check`

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all “i” series platforms.

disable slb gogo-mode service-check

```
disable slb gogo-mode <port number> service-check [all | ftp | http | nntp
| pop3 | smtp | telnet | <TCP port number>]
```

Description

Disables layer 7 service check to this GoGo mode group.

Syntax Description

port number	Specifies the GoGo mode master port.
all	Specifies all service checks.
ftp	Specifies the FTP service check.
http	Specifies the HTTP service check.
nntp	Specifies the NNTP service check.
pop3	Specifies the POP3 service check.
smtp	Specifies the SMTP service check.
telnet	Specifies the telnet service check.
TCP port number	Specifies the TCP port, instead of the service.

Default

GoGo mode service check is disabled by default.

Usage Guidelines

To enable service-check for a GoGo mode group, use the following command:

```
enable slb gogo-mode <port number> service-check
```

Example

The following command disables GoGo mode FTP service-check for the group with port 29 as the master port:

```
disable slb gogo-mode 29 service-check ftp
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all “i” series platforms.

disable slb gogo-mode tcp-port-check

```
disable slb gogo-mode <port number> tcp-port-check [all | ftp | http |
https | imap4 | ldap | nntp | pop3 | smtp | socks | telnet | tftp | web |
www | <TCP port number>]
```

Description

Disables layer 4 TCP-port-check to this GoGo mode group.

Syntax Description

port number	Specifies the GoGo mode master port.
all	Specifies all TCP-port-checks.
ftp	Specifies the FTP TCP-port-check.
http	Specifies the HTTP TCP-port-check.
https	Specifies the HTTPS TCP-port-check.
imap4	Specifies the IMAP4 TCP-port-check.
ldap	Specifies the LDAP TCP-port-check.
nntp	Specifies the NNTP TCP-port-check.
pop3	Specifies the POP3 TCP-port-check.
smtp	Specifies the SMTP TCP-port-check.
socks	Specifies the SOCKS TCP-port-check.
telnet	Specifies the telnet TCP-port-check.
tftp	Specifies the TFTP TCP-port-check.
web	Specifies the Web TCP-port-check.
www	Specifies the www TCP-port-check.
TCP port number	Specifies the TCP port of the TCP-port-check.

Default

GoGo mode TCP-port-check is disabled by default.

Usage Guidelines

To enable TCP-port-check for a GoGo mode group, use the following command:

```
enable slb gogo-mode <port number> tcp-port-check
```

Example

The following command disables all GoGo mode TCP-port-checks for the group with port 29 as the master port:

```
disable slb gogo-mode 29 tcp-port-check all
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all “i” series platforms.

disable slb L4-port

```
disable slb L4-port [all | ftp | http | https | imap4 | ldap | nntp | pop3
| smtp | socks | telnet | tftp | web | www | <TCP or UDP port number>]
```

Description

Disables one or all SLB ports.

Syntax Description

all	Specifies all nodes.
ftp	Specifies an FTP node.
http	Specifies an HTTP node.
https	Specifies an HTTPS node.
imap4	Specifies an IMAP4 node.
ldap	Specifies an LDAP node.
nntp	Specifies an NNTP node.
pop3	Specifies a POP3 node.
smtp	Specifies an SMTP node.
socks	Specifies a SOCKS node.
telnet	Specifies a telnet node.
tftp	Specifies a TFTP node.
web	Specifies a Web node.
www	Specifies a www node
TCP or UDP port number	Specifies a TCP or UDP port for the node.

Default

N / A.

Usage Guidelines

To enable an SLB port, use the following command:

`enable slb L4-port`

Example

The following command disables SLB for FTP ports:

`disable slb L4-port ftp`

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

disable slb node

```
disable slb node [all | <ip address> : [ftp | http | https | imap4 | ldap |
nntp | pop3 | smtp | socks | telnet | tftp | web | www | <TCP or UDP port
number>]] {close-connections-now}
```

Description

Disables one or all nodes.

Syntax Description

all	Specifies all nodes.
ip address	Specifies an IP address.
ftp	Specifies an FTP node.
http	Specifies an HTTP node.
https	Specifies an HTTPS node.
imap4	Specifies an IMAP4 node.
ldap	Specifies an LDAP node.
nntp	Specifies an NNTP node.
pop3	Specifies a POP3 node.
smtp	Specifies an SMTP node.
socks	Specifies a SOCKS node.
telnet	Specifies a telnet node.
tftp	Specifies a TFTP node.
web	Specifies a Web node.
www	Specifies a www node
TCP or UDP port number	Specifies a TCP or UDP port for the node.
close-connections-now	Immediately closes all open connections.

Default

N/A.

Usage Guidelines

This command stops nodes from accepting new connections; existing connections are not closed unless you specify `close-connections-now`. SLB continues to function with other nodes.

If you disable all nodes in a pool, all virtual servers associated with that pool are effectively disabled.

To enable a node, use the following command:

`enable slb node`

Example

The following command disables all nodes and immediately closes all open connections:

```
disable slb node all close-connections-now
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

disable slb node ping-check

```
disable slb node [all | <ip address>] ping-check
```

Description

Disables layer 3 ping-check.

Syntax Description

all	Specifies all nodes.
ip address	Specifies the IP address of the node.

Default

Ping-check is disabled by default.

Usage Guidelines

Ping-check is automatically enabled when a node is added to a pool.

To enable ping-check on a node, use the following command:

`enable slb node ping-check`

Example

The following command disables all ping-checks:

`disable slb node all ping-check`

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

disable slb node tcp-port-check

```
disable slb node [all | <ip address> : [ftp | http | https | imap4 | ldap |
nntp | pop3 | smtp | socks | telnet | tftp | web | www | <TCP or UDP port
number>]] tcp-port-check
```

Description

Disables layer 4 TCP-port-checking.

Syntax Description

all	Specifies all nodes.
ip address	Specifies an IP address.
ftp	Specifies an FTP node.
http	Specifies an HTTP node.
https	Specifies an HTTPS node.
imap4	Specifies an IMAP4 node.
ldap	Specifies an LDAP node.
nntp	Specifies an NNTP node.
pop3	Specifies a POP3 node.
smtp	Specifies an SMTP node.
socks	Specifies a SOCKS node.
telnet	Specifies a telnet node.
tftp	Specifies a TFTP node.
web	Specifies a Web node.
www	Specifies a www node
TCP or UDP port number	Specifies a TCP or UDP port for the node.

Default

TCP-port-check is disabled by default.

Usage Guidelines

To enable TCP-port-check, use the following command:

```
enable slb node tcp-port-check
```

Example

The following command disables all TCP-port-checks:

```
disable slb node all tcp-port-check
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

disable slb proxy-client-persistence

```
disable slb proxy-client-persistence
```

Description

Disables proxy client persistence.

Syntax Description

This command has no arguments or variables.

Default

Proxy client persistence is disabled by default.

Usage Guidelines

To enable proxy client persistence, use the following command:

`enable slb proxy-client-persistence`

Example

The following command disables proxy client persistence:

`disable slb proxy-client-persistence`

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

disable slb vip

```
disable slb vip [all | <vip name> | ipaddress <ip address> : [ftp | http |
https | imap4 | ldap | nntp | pop3 | smtp | socks | telnet | tftp | web |
www | <TCP or UDP port number>]] {close-connections-now}
```

Description

Disables one or all virtual servers.

Syntax Description

all	Specifies all virtual servers.
vip name	Specifies a virtual server.
ip address	Specifies an IP address.
ftp	Specifies an FTP virtual server.
http	Specifies an HTTP virtual server.
https	Specifies an HTTPS virtual server.
imap4	Specifies an IMAP4 virtual server.
ldap	Specifies an LDAP virtual server.
nntp	Specifies an NNTP virtual server.
pop3	Specifies a POP3 virtual server.
smtp	Specifies an SMTP virtual server.
socks	Specifies a SOCKS virtual server.
telnet	Specifies a telnet virtual server.
tftp	Specifies a TFTP virtual server.
web	Specifies a Web virtual server.
www	Specifies a www virtual server
TCP or UDP port number	Specifies a TCP or UDP port for the virtual server.
close-connections-now	Immediately closes all open connections.

Default

SLB is disabled by default.

Usage Guidelines

When disabled, no new connections are allowed to the real servers. If `close-connections-now` is specified, all existing connections are immediately closed.

To enable a virtual server, use the following command:

`enable slb vip`

Example

The following command disables the virtual server “ftp_vip” and closes all open connections:

```
disable slb vip ftp_vip close-connections-now
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

disable slb vip client-persistence

```
disable slb vip [all | <vip name>] client-persistence
```

Description

Disables client persistence.

Syntax Description

all	Specifies all virtual servers.
vip name	Specifies a virtual server.

Default

Client persistence is disabled by default.

Usage Guidelines

To enable client persistence, use the following command:

`enable slb vip client-persistence`

Example

The following command disables client persistence for the virtual server “ftp_vip”:

`disable slb vip ftp_vip client-persistence`

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

disable slb vip service-check

```
disable slb vip [all | <vip name>] service-check
```

Description

Disables layer 7 service-check.

Syntax Description

all	Specifies all virtual servers.
vip name	Specifies a virtual server.

Default

Service-check is disabled by default.

Usage Guidelines

To enable service-check, use the following command:

```
enable slb vip service-check
```

Example

The following command disables service-check for the virtual server “ftp_vip”:

```
disable slb vip ftp_vip service-check
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

disable slb vip sticky-persistence

```
disable slb vip [all | <vip name>] sticky-persistence
```

Description

Disables sticky persistence.

Syntax Description

all	Specifies all virtual servers.
vip name	Specifies a virtual server.

Default

Sticky persistence is disabled by default.

Usage Guidelines

To enable sticky persistence, use the following command:

`enable slb vip sticky-persistence`

Example

The following command disables sticky persistence for the virtual server “ftp_vip”:

`disable slb vip ftp_vip sticky-persistence`

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

disable slb vip svcdown-reset

```
disable slb vip [all | <vip name>] svcdown-reset
```

Description

Disables svcdown-reset.

Syntax Description

all	Specifies all virtual servers.
vip name	Specifies a virtual server.

Default

The svcdown-reset feature is disabled by default.

Usage Guidelines

To enable svcdown-reset, use the following command:

```
enable slb vip svcdown-reset
```

Example

The following command disables svcdown-reset for the virtual server “ftp_vip”:

```
disable slb vip ftp_vip svcdown-reset
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

enable flow-redirect

```
enable flow-redirect [all | <flow redirect>]
```

Description

Enables flow redirect.

Syntax Description

all	Specifies all flow policies.
flow redirect	Specifies a single flow redirect policy.

Default

The default parameter is all.

Flow redirection is disabled by default.

Usage Guidelines

When you create a new flow redirect policy, flow redirect is automatically enabled.

To disable flow redirect, use the following command:

```
disable flow-redirect
```

Example

The following command enables flow redirect for all flow policies:

```
enable flow-redirect all
```

History

This command was available in ExtremeWare 6.1.4.

Platform Availability

This command is available on all "*i*" series platforms.

enable slb

```
enable slb
```

Description

Enables SLB processing.

Syntax Description

This command has no arguments or variables.

Default

SLB is disabled by default.

Usage Guidelines

This command activates the following functions for transparent, translational, and port translation modes:

- Exporting of VIP routes or proxy ARP for VIP addresses.
- Processing of VIP lookup and connection setup.
- Establishing communication with redundant SLB switches.
- Positively responding to MIB, 3DNS, and SeeIT requests.

Before you enable SLB, enable IP forwarding on the associated VLANs.

Example

The following command enables SLB:

```
enable slb
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

enable slb 3dns

```
enable slb 3dns iquery-client
```

Description

Enables 3DNS support.

Syntax Description

This command has no arguments or variables.

Default

3DNS is disabled by default.

Usage Guidelines

The following 3DNS global balance modes are supported:

- completion
- rate
- global_availability
- leastcon
- null
- packet_rate
- random
- ration
- rr
- return_to_dns

To disable 3DNS, use the following command:

```
disable slb 3dns iquery-client
```

Example

The following command enables 3DNS:

```
enable slb 3dns iquery-client
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

enable slb failover

```
enable slb failover
```

Description

Enables SLB failover.

Syntax Description

This command has no arguments or variables.

Default

Failover is disabled by default.

Usage Guidelines

When SLB failover is enabled, the primary SLB switch automatically resumes primary status when it becomes active.

Before you enable SLB failover, configure your switches using the following command:

```
config slb failover unit
```

To disable SLB failover, use the following command:

```
disable slb failover
```

Example

The following command enables SLB failover:

```
enable slb failover
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

enable slb failover manual-failback

```
enable slb failover manual-failback
```

Description

Enables manual failback.

Syntax Description

This command has no arguments or variables.

Default

Manual failback is disabled by default.

Usage Guidelines

When manual failback is enabled, the primary SLB switch does not automatically resume primary status until you use the following command:

```
config slb failover failback-now
```

To disable manual failback, use the following command:

```
disable slb failover manual-failback
```

Example

The following command enables manual failback:

```
enable slb failover manual-failback
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

enable slb failover ping-check

```
enable slb failover ping-check
```

Description

Enables ping-check.

Syntax Description

This command has no arguments or variables.

Default

Ping-check is disabled by default.

Usage Guidelines

To disable ping-check, use the following command:

```
disable slb failover ping-check
```

Example

The following command enables ping-check:

```
enable slb failover ping-check
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

enable slb global synguard

```
enable slb global synguard
```

Description

Enables the TCP SYN-guard feature.

Syntax Description

This command has no arguments or variables.

Default

SYN-guard is disabled by default.

Usage Guidelines

To disable SYN-guard, use the following command:

```
disable slb global synguard
```

Example

The following command enables SYN-guard:

```
enable slb global synguard
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

enable slb gogo-mode

```
enable slb gogo-mode <port number> grouping <port list>
```

Description

Enables GoGo mode processing for a group of ports.

Syntax Description

port number	Specifies the GoGo mode master port.
port list	Specifies a range or list of ports assigned to the group.

Default

GoGo mode is disabled by default.

Usage Guidelines

To disable GoGo mode, use the following command:

```
disable slb gogo-mode
```

Example

The following command enables GoGo mode for the group containing ports 15, 17, 19-23, and 25-30 with port 29 as the master port:

```
enable slb gogo-mode 29 grouping 15, 17, 19 - 23, 25 - 30
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

enable slb gogo-mode ping-check

```
enable slb gogo-mode <port number> ping-check <ip address>
```

Description

Enables layer-3 ping-check for the GoGo mode group.

Syntax Description

port number	Specifies the GoGo mode master port.
ip address	Specifies an IP address to be pinged.

Default

GoGo mode ping check is disabled by default.

Usage Guidelines

GoGo mode ping-check sends a ping from each physical port in the GoGo mode grouping to the configured IP address.

If you do not specify an IP address, GoGo mode ping-check uses the previously configured IP address.

You must enable GoGo mode for the group before you enable ping-check. To enable GoGo mode, use the following command:

```
enable slb gogo-mode
```

To disable ping-check for a GoGo mode group, use the following command:

```
disable slb gogo-mode <port number> ping-check
```

Example

The following command enables GoGo mode ping-check for the group with port 29 as the master port to IP address 10.10.200.3:

```
enable slb gogo-mode 29 ping-check 10.10.200.3
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all “i” series platforms.

enable slb gogo-mode service-check

```
enable slb gogo-mode <port number> service-check [all | ftp | http | nntp |
pop3 | smtp | telnet | <TCP port number>]
```

Description

Enables layer 7 service checking for the GoGo mode group.

Syntax Description

port number	Specifies the GoGo mode master port.
all	Specifies all service checks.
ftp	Specifies the FTP service check.
http	Specifies the HTTP service check.
nntp	Specifies the NNTP service check.
pop3	Specifies the POP3 service check.
smtp	Specifies the SMTP service check.
telnet	Specifies the telnet service check.
TCP port number	Specifies the TCP port, instead of the service.

Default

GoGo mode service check is disabled by default.

Usage Guidelines

To disable service-check for a GoGo mode group, use the following command:

```
disable slb gogo-mode <port number> service-check
```

Example

The following command enables GoGo mode FTP service-check for the group with port 29 as the master port:

```
enable slb gogo-mode 29 service-check ftp
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all “i” series platforms.

enable slb gogo-mode tcp-port-check

```
enable slb gogo-mode <port number> tcp-port-check [all | ftp | http | https
| imap4 | ldap | nntp | pop3 | smtp | socks | telnet | tftp | web | www |
<TCP port number>]
```

Description

Enables layer 4 TCP-port-check for the GoGo mode group.

Syntax Description

port number	Specifies the GoGo mode master port.
all	Specifies all TCP-port-checks.
ftp	Specifies the FTP TCP-port-check.
http	Specifies the HTTP TCP-port-check.
https	Specifies the HTTPS TCP-port-check.
imap4	Specifies the IMAP4 TCP-port-check.
ldap	Specifies the LDAP TCP-port-check.
nntp	Specifies the NNTP TCP-port-check.
pop3	Specifies the POP3 TCP-port-check.
smtp	Specifies the SMTP TCP-port-check.
socks	Specifies the SOCKS TCP-port-check.
telnet	Specifies the telnet TCP-port-check.
tftp	Specifies the TFTP TCP-port-check.
web	Specifies the Web TCP-port-check.
www	Specifies the www TCP-port-check.
TCP port number	Specifies the TCP port of the TCP-port-check.

Default

GoGo mode TCP-port-check is disabled by default.

Usage Guidelines

To disable TCP-port-check for a GoGo mode group, use the following command:

```
disable slb gogo-mode <port number> tcp-port-check
```

Example

The following command enables all GoGo mode TCP-port-checks for the group with port 29 as the master port:

```
enable slb gogo-mode 29 tcp-port-check all
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all “i” series platforms.

enable slb L4-port

```
enable slb L4-port [ftp | http | https | imap4 | ldap | nntp | pop3 | smtp
| socks | telnet | tftp | web | www | <TCP or UDP port number>]
```

Description

Enables an SLB port.

Syntax Description

ftp	Specifies an FTP node.
http	Specifies an HTTP node.
https	Specifies an HTTPS node.
imap4	Specifies an IMAP4 node.
ldap	Specifies an LDAP node.
nntp	Specifies an NNTP node.
pop3	Specifies a POP3 node.
smtp	Specifies an SMTP node.
socks	Specifies a SOCKS node.
telnet	Specifies a telnet node.
tftp	Specifies a TFTP node.
web	Specifies a Web node.
www	Specifies a www node
TCP or UDP port number	Specifies a TCP or UDP port for the node.

Default

Layer 4 ports are enabled by default.

Usage Guidelines

To disable a layer 4 port, use the following command:

```
disable slb L4-port
```

Example

The following command enables SLB for FTP ports:

```
enable slb L4-port ftp
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

enable slb node

```
enable slb node [all | <ip address> : [ftp | http | https | imap4 | ldap |
nntp | pop3 | smtp | socks | telnet | tftp | web | www | <TCP or UDP port
number>]]
```

Description

Enables one or all nodes.

Syntax Description

all	Specifies all nodes.
ip address	Specifies an IP address.
ftp	Specifies an FTP node.
http	Specifies an HTTP node.
https	Specifies an HTTPS node.
imap4	Specifies an IMAP4 node.
ldap	Specifies an LDAP node.
nntp	Specifies an NNTP node.
pop3	Specifies a POP3 node.
smtp	Specifies an SMTP node.
socks	Specifies a SOCKS node.
telnet	Specifies a telnet node.
tftp	Specifies a TFTP node.
web	Specifies a Web node.
www	Specifies a www node
TCP or UDP port number	Specifies a TCP or UDP port for the node.

Default

Nodes are enabled by default.

Usage Guidelines

This command allows nodes to accept new connections.

To disable a node, use the following command:

```
disable slb node
```

Example

The following command enables all nodes:

```
enable slb node all
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

enable slb node ping-check

```
enable slb node [all | <ip address>] ping-check
```

Description

Enables layer 3 ping-check.

Syntax Description

all	Specifies all nodes.
ip address	Specifies the IP address of the node.

Default

Ping-check is enabled by default.

Usage Guidelines

Ping-check is automatically enabled when a node is added to a pool.

To disable ping-check on a node, use the following command:

```
disable slb node ping-check
```

Example

The following command enables all ping-checks:

```
enable slb node all ping-check
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

enable slb node tcp-port-check

```
enable slb node [all | <ip address> : [ftp | http | https | imap4 | ldap |
nntp | pop3 | smtp | socks | telnet | tftp | web | www | <TCP or UDP port
number>]] tcp-port-check
```

Description

Enables layer 4 TCP-port-check.

Syntax Description

all	Specifies all nodes.
ip address	Specifies an IP address.
ftp	Specifies an FTP node.
http	Specifies an HTTP node.
https	Specifies an HTTPS node.
imap4	Specifies an IMAP4 node.
ldap	Specifies an LDAP node.
nntp	Specifies an NNTP node.
pop3	Specifies a POP3 node.
smtp	Specifies an SMTP node.
socks	Specifies a SOCKS node.
telnet	Specifies a telnet node.
tftp	Specifies a TFTP node.
web	Specifies a Web node.
www	Specifies a www node
TCP or UDP port number	Specifies a TCP or UDP port for the node.

Default

TCP-port-check is disabled by default.

Usage Guidelines

To disable TCP-port-check, use the following command:

```
disable slb node tcp-port-check
```

Example

The following command enables all TCP-port-checks:

```
enable slb node all tcp-port-check
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

enable slb proxy-client-persistence

```
enable slb proxy-client-persistence
```

Description

Enables proxy client persistence.

Syntax Description

This command has no arguments or variables.

Default

Proxy client persistence is disabled by default.

Usage Guidelines

To disable proxy client persistence, use the following command:

`disable slb proxy-client-persistence`

Example

The following command enables proxy client persistence:

`enable slb proxy-client-persistence`

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

enable slb vip

```
enable slb vip [all | <vip name> | ipaddress <ip address> : [ftp | http |
https | imap4 | ldap | nntp | pop3 | smtp | socks | telnet | tftp | web |
www | <TCP or UDP port number>]]
```

Description

Enables one or all virtual servers.

Syntax Description

all	Specifies all virtual servers.
vip name	Specifies a virtual server.
ip address	Specifies an IP address.
ftp	Specifies an FTP virtual server.
http	Specifies an HTTP virtual server.
https	Specifies an HTTPS virtual server.
imap4	Specifies an IMAP4 virtual server.
ldap	Specifies an LDAP virtual server.
nntp	Specifies an NNTP virtual server.
pop3	Specifies a POP3 virtual server.
smtp	Specifies an SMTP virtual server.
socks	Specifies a SOCKS virtual server.
telnet	Specifies a telnet virtual server.
tftp	Specifies a TFTP virtual server.
web	Specifies a Web virtual server.
www	Specifies a www virtual server
TCP or UDP port number	Specifies a TCP or UDP port for the virtual server.

Default

SLB is disabled by default.

Usage Guidelines

To disable a virtual server, use the following command:

```
disable slb vip
```

Example

The following command enables the virtual server “ftp_vip”:

```
enable slb vip ftp_vip
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

enable slb vip client-persistence

```
enable slb vip [all | <vip name>] client-persistence {netmask <netmask>}
```

Description

Enables client persistence.

Syntax Description

all	Specifies all virtual servers.
vip name	Specifies a virtual server.
netmask	Specifies a netmask.

Default

The default is disabled.

Usage Guidelines

To disable client persistence, use the following command:

```
disable slb vip client-persistence
```

Example

The following command enables client persistence for the virtual server “ftp_vip”:

```
enable slb vip ftp_vip client-persistence
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

enable slb vip service-check

```
enable slb vip [all | <vip name>] service-check
```

Description

Enables layer 7 service check.

Syntax Description

all	Specifies all virtual servers.
vip name	Specifies a virtual server.

Default

Service-check is disabled by default.

Usage Guidelines

The service checks are based on the following information:

- If a service check is already configured, then it will use the configured service-checking information.
- If a service-check is configured for a TCP port number (instead of for a service), ExtremeWare assigns the service based on the port number (if the port number is well known) and uses the global default parameters.

To disable service-check, use the following command:

```
disable slb vip service-check
```

Example

The following command enables service-check for the virtual server “ftp_vip”:

```
enable slb vip ftp_vip service-check
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

enable slb vip sticky-persistence

```
config slb vip [all | <vip name>] sticky-persistence {netmask <netmask>}
```

Description

Enables the sticky persistence feature and specifies the client address mask.

Syntax Description

all	Specifies all virtual servers.
vip name	Specifies a virtual server.
netmask	Specifies a netmask.

Default

Sticky persistence is disabled by default.

Usage Guidelines

To disable sticky persistence, use the following command:

```
disable slb vip sticky-persistence
```

Example

The following command enables sticky persistence for the virtual server “ftp_vip”:

```
enable slb vip ftp_vip sticky-persistence
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

enable slb vip svcdown-reset

```
enable slb vip [all | <vipname>] svcdown-reset
```

Description

Enables svcdown-reset.

Syntax Description

all	Specifies all virtual servers.
vip name	Specifies a virtual server.

Default

The svcdown-reset feature is disabled by default.

Usage Guidelines

The svcdown-reset feature configures the switch to send TCP RST packets to both the clients and the virtual server if the virtual server fails a health-check.

To disable svcdown-reset, use the following command:

```
disable slb vip svcdown-reset
```

Example

The following command enables svcdown-reset for the virtual server “ftp_vip”:

```
enable slb vip ftp_vip svcdown-reset
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

show flow-redirect

```
show flow-redirect <flow redirect>
```

Description

Displays the current flow redirect configuration and statistics.

Syntax Description

flow redirect	Specifies a flow redirect policy.

Default

N / A.

Usage Guidelines

If you do not specify a flow redirect policy, configuration and statistics for all flow redirect policies are displayed.

Following are the fields displayed:

Service Check Timer Settings:	Displays the frequency and timeout settings for the flow-redirect ping-check, TCP-port-check, and service-check.
Flow IPSA Mode	Displays the IP source address mode: • Enumeration Mode—The default mode, used for network masks from /32 to /20. • Subnet Mode—Used for network masks from /19 to /1. The mode is selected automatically when you specify a network mask.
Proto:	Displays the flow type. • any—Forwards any traffic over any IP port. • tcp—Forwards TCP traffic over a single IP port. • tup—Forwards both TCP and UDP traffic over a single IP port. • udp—Forwards UDP traffic over a single IP port.
Dest:	Displays the desitnation IP address.
L4-port:	Displays the port number.
Enabled:	Displays status of flow-redirect. • Yes—Flow redirect is enabled. • No—Flow redirect is not enabled.
Source:	• Displays the source IP address.
# Servers Up:	Displays the number of next hops up over the number of next hops configured.

Service Checking:	Displays the configured service check type. • ftp • http • L4-port • nntp • ping • pop3 • smtp • telnet
IP Address	Displays the IP address of the next hop.
State	Displays the status of the next hop, either up or down.
Flow Info	Displays hardware mapping information.

Example

The following command displays the current flow redirect configuration and statistics for the flow policy “flow1”:

```
show flow-redirect flow1
```

Following is the output from this command:

```
Service Check Timer Settings:
  Ping-check     Frequency:  10  Timeout:  30
  TCP-Port-check Frequency:  10  Timeout:  30
  Service-check  Frequency:  60  Timeout: 180

Flow IPSA Mode: Enumeration Mode

http1
        Proto:tcp  Dest:      0.0.0.0/ 0  L4-Port:    80   Enabled: yes
                   Source:    0.0.0.0/ 0              # Servers Up: 0/1
        Service Checking: ping
        IP Address      State   Flow Info
        24.3.89.145     Down    0000
```

History

This command was available in ExtremeWare 6.1.4.

Platform Availability

This command is available on all “*i*” series platforms.

show slb 3dns members

```
show slb 3dns members
```

Description

Displays the current connection information between the switch and the 3DNS querier.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the current 3DNS information:

```
show slb 3dns members
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

show slb connections

```
show slb connections [ipaddress <ip address>: [ftp | http | https | imap4 |
ldap | nntp | pop3 | smtp | socks | telnet | tftp | web | www | <TCP or UDP
port number>] | vip <vip name>]
```

Description

Displays information on current connections.

Syntax Description

ip address	Specifies an IP address.
vip name	Specifies a virtual server.
ftp	Specifies an FTP port.
http	Specifies an HTTP port.
https	Specifies an HTTPS port.
imap4	Specifies an IMAP4 port.
ldap	Specifies an LDAP port.
nntp	Specifies an NNTP port.
pop3	Specifies a POP3 port.
smtp	Specifies an SMTP port.
socks	Specifies a SOCKS port.
telnet	Specifies a telnet port.
tftp	Specifies a TFTP port.
web	Specifies a Web port.
www	Specifies a www port
TCP or UDP port number	Specifies a TCP or UDP port.

Default

N/A.

Usage Guidelines

You can specify a client, virtual server, or node. If you do not specify a virtual server or IP address, information on all connections is displayed. An IP address of 0.0.0.0:0 is a wildcard.

Example

The following command displays the current connection information for all connections:

```
show slb connections
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

show slb esrp

```
show slb esrp
```

Description

Displays SLB configuration for ESRP.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the current ESRP configuration:

```
show slb esrp
```

Following is the output from this command:

```
VLAN Name  SLB Unit Status    SLB Unit(s)
servers    Standby            1
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

show slb failover

```
show slb failover
```

Description

Displays SLB failover configuration and status.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The `show slb global` command also displays SLB failover configuration and status.

Example

The following command displays the current SLB failover configuration and status:

```
show slb failover
```

Following is the output from this command:

```
SLB Failover Configuration:
 Failover: Enabled
 Local unit ID: 1
 Local IP address: 10.1.1.1
 Remote IP address: 10.1.1.2
 TCP port number: 1028
 Remote Alive frequency: 1
 Remote Dead frequency: 2
 Keepalive Timeout: 3
 Ping check: Disabled
 Ping check IP address: 0.0.0.0
 Ping frequency: 1
 Ping timeout: 3
 Manual failback: Disabled
#
#
SLB Failover Status: Running
 Units active in local SLB: 2
 Units active in or
  requested by remote SLB: None
 Send connection: Down
 Receive connection: Down
 Ping check: Not Running
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

show slb global

```
show slb global
```

Description

Displays the current SLB global configuration information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Displays the following:

- Global enable/disable mode
- Global modes
- Default settings for health checker
- Failover configuration

Example

The following command displays the current SLB global configuration information:

```
show slb global
```

Following is the output from this command:

```
SLB: Enabled
SynGuard: Disabled
3DNS IQuery Support Status: Disabled
SLB persist-level: same-vip-same-port
SLB persistence-method: per-session
SLB pre-allocated connection-block size: 10000
SLB connection timeout: 1
SLB persistence on client proxies: Disabled
Proxy Client Persistence entries:
No. of Proxy Client Persistence entries: 0
#
#
Health Check Defaults:
  Ping-check    Frequency: 10  Timeout: 30
  Port-check    Frequency: 30  Timeout: 90
  Service-check Frequency: 60  Timeout: 180
  HTTPURL: "/"
      Match String: (any-content)
  FTP User: anonymous
      Password: (not shown)
  TelnetUser: anonymous
```

```
      Password: (not shown)
  SMTPDomain: "mydomain.com"
      NNTP Newsgroup: "ebusiness"
      User: anonymous
      Password: (not shown)
  POP3User: anonymous
      Password: (not shown)
#
#
SLB Failover Configuration:
Failover: Enabled
 Local unit ID: 1
 Local IP address: 10.1.1.1
 Remote IP address: 10.1.1.2
 TCP port number: 1028
 Remote Alive frequency: 1
 Remote Dead frequency: 2
 Keepalive Timeout: 3
 Ping check: Disabled
 Ping check IP address: 0.0.0.0
 Ping frequency: 1
 Ping timeout: 3
 Manual failback: Disabled
#
#
SLB Failover Status: Running
 Units active in local SLB: 2
 Units active in or
  requested by remote SLB: None
 Send connection: Down
 Receive connection: Down
 Ping check: Not Running
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

show slb gogo-mode

```
show slb gogo-mode <port number> {configuration}
```

Description

Displays GoGo mode ping-check, TCP-port-check, and service-check status.

Syntax Description

port number	Specifies the GoGo mode master port.
configuration	Displays configuration instead of status.

Default

N/A.

Usage Guidelines

If you do not specify a master port, status for all GoGo mode groups with health checks configured is displayed.

Example

The following command displays the current GoGo mode health check configuration for the group with port 29 as the master port:

```
show slb gogo-mode 29 configuration
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all “i” series platforms.

show slb L4-port

```
show slb L4-port [ftp | http | https | imap4 | ldap | nntp | pop3 | smtp |
socks | telnet | tftp | web | www | <TCP or UDP port number>]
```

Description

Displays the SLB configuration for the active layer 4 ports.

Syntax Description

ftp	Specifies an FTP port.
http	Specifies an HTTP port.
https	Specifies an HTTPS port.
imap4	Specifies an IMAP4 port.
ldap	Specifies an LDAP port.
nntp	Specifies an NNTP port.
pop3	Specifies a POP3 port.
smtp	Specifies an SMTP port.
socks	Specifies a SOCKS port.
telnet	Specifies a telnet port.
tftp	Specifies a TFTP port.
web	Specifies a Web port.
www	Specifies a www port
TCP or UDP port number	Specifies a TCP or UDP port.

Default

N / A.

Usage Guidelines

If you do not specify a port, configuration and status for all layer 4 ports is displayed.

Example

The following command displays the current layer 4 port configuration:

```
show slb L4-port
```

Following is the output from this command:

```
Port:    80  Enabled  TCP idle timeout (treaper): 600  UDP idle timeout: 600
```

History

This command was first available in ExtremeWare 6.1.

show slb node

```
show slb node {<ip address> [ftp | http | https | imap4 | ldap | nntp |
pop3 | smtp | socks | telnet | tftp | web | www | <TCP or UDP port
number>]}
```

Description

Displays node configuration and status.

Syntax Description

ip address	Specifies an IP address.
ftp	Specifies an FTP node.
http	Specifies an HTTP node.
https	Specifies an HTTPS node.
imap4	Specifies an IMAP4 node.
ldap	Specifies an LDAP node.
nntp	Specifies an NNTP node.
pop3	Specifies a POP3 node.
smtp	Specifies an SMTP node.
socks	Specifies a SOCKS node.
telnet	Specifies a telnet node.
tftp	Specifies a TFTP node.
web	Specifies a Web node.
www	Specifies a www node
TCP or UDP port number	Specifies a TCP or UDP port for the node.

Default

N/A.

Usage Guidelines

If you do not specify a node, status for all nodes is displayed.

Example

The following command displays the current node configuration and statistics for all nodes:

```
show slb node
```

Following is the output from this command:

```
                    IP     Freq/        TCP/UDP      Frequency/Max
Node IP Address    Flags   Timeout    Port    Flags   Timeout#PoolsConns
1.111.1.1          E--H--  10/30      80      E---    30/90 2(no limit)
1.111.1.2          E--H--  10/30      80      E---    30/90 2(no limit)
1.111.1.3          E--H--  10/30      80      E---    30/90 2(no limit)

Flags: E - Enable,   U - Up,   R - IP Route Up,   H - Health check enabled,
       P - Health check passed,   ! - VLAN not configured with "slb-type"
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

show slb persistence

```
show slb persistence
```

Description

Displays persistence status of existing clients.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the current persistence status:

```
show slb persistence
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

show slb pool

```
show slb pool <pool name>
```

Description

Displays the current SLB pool configuration and status.

Syntax Description

pool name	Specifies a pool.

Default

N / A.

Usage Guidelines

If you do not specify a pool, configuration and status for all pools is displayed.

Example

The following command displays the current pool configuration and statistics for all pools, currently “rr_pool” and “ratio_pool”:

```
show slb pool
```

Following is the output from this command:

```
Name                        IP                  TCP/UDP    Ratio/
      IP                   Flags      Port      Flags      Priority
-------------------------------------------------------------------------
rr_pool               # VIPs sharing: 1  Load Bal. Method: Round Robin
      1.111.1.1         E--H--     80      E---
      1.111.1.2         E--H--     80      E---
      1.111.1.3         E--H--     80      E---
ratio_pool            # VIPs sharing: 1  Load Bal. Method: Ratio
      1.111.1.3         E--H--     80      E---      3
      1.111.1.2         E--H--     80      E---      2
      1.111.1.1         E--H--     80      E---      1

Flags: E - Enable,   U - Up,   R - IP Route Up,   H - Health check enabled,
       P - Health check passed,   ! - VLAN not configured with "slb-type"
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

Platform Availability

This command is available on all “i” series platforms.

show slb stats

```
show slb stats [pool <pool name> | vip <vip name>]
```

Description

Displays the current SLB pool connection status.

Syntax Description

pool name	Specifies a pool.
vip name	Specifies a virtual server.

Default

N/A.

Usage Guidelines

If you specify `pool` but do not specify a specific pool, status for all pools is displayed.

If you specify `vip` but do not specify a specifuc virtual server, status for all virtual servers is displayed.

If you do not specify a pool or virtual server, status for all pools and virtual servers is displayed.

Example

The following command displays the current pool connection status for all pools:

```
show slb stats pool
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

show slb vip

```
show slb vip [<vip name> | ipaddress <ip address> : [ftp | http | https |
imap4 | ldap | nntp | pop3 | smtp | socks | telnet | tftp | web | www |
<TCP or UDP port number>]] {detail}
```

Description

Displays the current virtual server configuration and statistics.

Syntax Description

vip name	Specifies a virtual server.
ip address	Specifies an IP address.
ftp	Specifies an FTP virtual server.
http	Specifies an HTTP virtual server.
https	Specifies an HTTPS virtual server.
imap4	Specifies an IMAP4 virtual server.
ldap	Specifies an LDAP virtual server.
nntp	Specifies an NNTP virtual server.
pop3	Specifies a POP3 virtual server.
smtp	Specifies an SMTP virtual server.
socks	Specifies a SOCKS virtual server.
telnet	Specifies a telnet virtual server.
tftp	Specifies a TFTP virtual server.
web	Specifies a Web virtual server.
www	Specifies a www virtual server
TCP or UDP port number	Specifies a TCP or UDP port for the virtual server.

Default

N/A.

Usage Guidelines

If you do not specify a virtual server or IP address, information on all virtual servers is displayed.

Example

The following command displays the current virtual server configuration and statistics for all virtual servers, currently “ratio_vip” and “rr_vip”:

```
show slb vip
```

Following is the output from this command:

```
                                        Unit   Export     # Servers
Name              IP Address         Port -- Mode --  FlagsPool     Up/Defined
ratio_vip         4.1.1.100          80   1  TL  SR  EUA-----ratio_po0/3
rr_vip            10.1.1.10          80   1  TP  PA  EUA----!rr_pool0/3

Modes: TP - Transparent,  TL - Translational,  PT - Port Translational
Automatically Exported via: PA - Proxy Arp, HR - Host Route, SR - Subnet Route
Flags: E - Enable,  U - Up,  A - Active Unit,  H - Health-Check Enabled,
       P - Persistence,  S - Sticky,  R - SvcDown-Reset,
       ! - VLAN has not been configured with "slb-type"
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

unconfig slb all

```
unconfig slb all
```

Description

Resets SLB global defaults and clears the SLB configuration.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command does not delete nodes, pools, or virtual servers. To delete all nodes and pools, use the following command:

```
delete slb pool all
```

To delete all virtual servers, use the following command:

```
delete slb vip all
```

Example

The following command resets SLB global defaults and clears the SLB configuration:

```
unconfig slb all
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

unconfig slb gogo-mode health-check

```
unconfig slb gogo-mode <port number> health-check
```

Description

Disables and deletes all the ping-check, TCP-port-check, and service-check configurations for this GoGo mode group.

Syntax Description

port number	Specifies the GoGo mode master port.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes all health-check configurations for the GoGo mode group with port 29 as the master port:

```
unconfig slb gogo-mode 29 health-check
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all “i” series platforms.

unconfig slb gogo-mode service-check

```
unconfig slb gogo-mode <port number> service-check [all | ftp | http | nntp
| pop3 | smtp | telnet | <TCP port number>]
```

Description

Disables and deletes the GoGo mode service-check configuration.

Syntax Description

port number	Specifies the GoGo mode master port.
all	Specifies all service checks.
ftp	Specifies the FTP service check.
http	Specifies the HTTP service check.
nntp	Specifies the NNTP service check.
pop3	Specifies the POP3 service check.
smtp	Specifies the SMTP service check.
telnet	Specifies the telnet service check.
TCP port number	Specifies the TCP port, instead of the service.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables and deletes all the FTP service-check configuration for the GoGo mode group with port 29 as the master port:

```
unconfig slb gogo-mode 29 service-check ftp
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all “i” series platforms.

unconfig slb vip service-check

```
unconfig slb vip [all | <vip name>] service-check
```

Description

Disables and deletes the service check configuration.

Syntax Description

all	Specifies all virtual servers.
vip name	Specifies a virtual server.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables and deletes the FTP service-check configurations for the virtual server “ftp_vip:

```
unconfig slb vip ftp_vip service-check
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “i” series platforms.

11 EAPS Commands

This chapter describes the following commands:

- Commands for configuring and monitoring Ethernet Automatic Protection Switching (EAPS)

To use EAPS, you must enable EDP on the switch.

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. 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.

To take advantage of the Spatial Reuse technology and broaden the use of the ring's bandwidth, EAPS supports multiple EAPS domains running on the ring at the same time.

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

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. If the master node detects a break in the ring, it unblocks its secondary port and allows data traffic to be transmitted and received through it.

EAPS 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.

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

- Failed response to a periodic health-check packet on the control VLAN
- "Link down" trap message sent by a transit node on the control VLAN

When the master node detects a failure, it declares a "failed" state and opens its logically blocked secondary port on all the protected VLANs. 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.

config eaps add control vlan

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

Description

Adds the specified control VLAN to the specified EAPS domain.

Syntax Description

name	Specifies the name of an EAPS domain.
vlan_name	Specifies the name of the control VLAN.

Default

N / A.

Usage Guidelines

You must configure one *control* VLAN for each EAPS domain. The control VLAN is used only to send and receive EAPS messages.

The VLAN that will act as the control VLAN must be configured as follows:

- The VLAN must NOT be assigned an IP address, to avoid loops in the network.
- Only ring ports may be added as members of the control VLAN.
- The ring ports of the control VLAN must be tagged. This ensures that EAPS control VLAN traffic is serviced before any other traffic and that control VLAN messages reach their intended destinations.
- The control VLAN must be assigned a QoS profile of QP8 with the QoS profile priority setting `HighHi`.

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

Example

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

```
config eaps eaps_1 add control vlan keys
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the “*i*” series platforms.

config eaps add protect vlan

```
config eaps <name> add protect vlan <vlan_name>
```

Description

Adds the specified protected VLAN to the specified EAPS domain.

Syntax Description

name	Specifies the name of an EAPS domain.
vlan_name	Specifies the name of the protected VLAN.

Default

N/A.

Usage Guidelines

You must configure one or more *protected* VLANs for each EAPS domain. The protected VLANs are the data-carrying VLANs.

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). As long as the ring is complete, the master node blocks the protected VLANs on its secondary port.

Example

The following command adds the protected VLAN "orchid" to the EAPS domain "eaps_1":

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

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config eaps delete control vlan

```
config eaps <name> delete control vlan <vlan_name>
```

Description

Deletes the specified control VLAN from the specified EAPS domain.

Syntax Description

name	Specifies the name of an EAPS domain.
vlan_name	Specifies the name of the control VLAN.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the control VLAN “keys” from the EAPS domain “eaps_1”:

```
config eaps eaps_1 delete control vlan keys
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the “*i*” series platforms.

config eaps delete protect vlan

```
config eaps <name> delete protect vlan <vlan_name>
```

Description

Deletes the specified protected VLAN from the specified EAPS domain.

Syntax Description

name	Specifies the name of an EAPS domain.
vlan_name	Specifies the name of the protected VLAN.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the protected VLAN “orchid” from the EAPS domain “eaps_1”:

```
config eaps eaps_1 delete protect vlan orchid
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the “*i*” series platforms.

config eaps failtime

```
config eaps <name> failtime <seconds>
```

Description

Configures the value of the fail timer the master node used for the EAPS health-check packet.

Syntax Description

name	Specifies the name of an EAPS domain.
seconds	Specifies the number of seconds to wait between transmission of the health-check packets on the control VLAN. Must be greater than 0.

Default

Default is 3 seconds.

Usage Guidelines

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.

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.

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.

Example

The following command configures the failtime value for the EAPS domain "eaps_1" to 10 seconds:

```
config eaps eaps_1 failtime 10
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config eaps hellotime

```
config eaps <name> hellotime <seconds>
```

Description

Configures the value of the hello timer the master node used for the EAPS health-check packet.

Syntax Description

name	Specifies the name of an EAPS domain.
seconds	Specifies the number of seconds to wait between transmission of the health-check packets on the control VLAN. Must be greater than 0.

Default

Default is 1 second.

Usage Guidelines

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. 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.

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.

Example

The following command configures the hellotime value for the EAPS domain “eaps_1” to 2 seconds:

```
config eaps eaps_1 hellotime 2
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the “*i*” series platforms.

config eaps mode

```
config eaps <name> mode [master | transit]
```

Description

Configures the switch as either the EAPS master node or as an EAPS transit node for the specified domain.

Syntax Description

name	Specifies the name of an EAPS domain.
master	Specifies that this switch should be the master node for the named EAPS domain.
transit	Specifies that this switch should be the transit node for the named EAPS domain.

Default

N/A.

Usage Guidelines

None.

Example

The following command identifies this switch as the master node for the domain named eaps_1:

```
config eaps eaps_1 master
```

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

```
config eaps eaps_1 transit
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config eaps port

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

Description

Configures a node port as the primary or secondary port for the specified EAPS domain.

Syntax Description

name	Specifies the name of an EAPS domain.
primary	Specifies that the port is to be configured as the primary port.
secondary	Specifies that the port is to be configured as the secondary port.
port number	Specifies the port number.

Default

N / A.

Usage Guidelines

Each node on the ring connects through two ring ports. One port must be configured as the *primary* port; the other must be configured as the *secondary* port.

Example

The following command adds port 1 of the module installed in slot 8 of a BlackDiamond switch to the EAPS domain “eaps_1” as the primary port:

```
config eaps eaps_1 primary port 8:1
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the “*i*” series platforms.

config eaps name

```
config eaps <old_name> name <new_name>
```

Description

Renames an existing EAPS domain.

Syntax Description

old_name	Specifies the current name of an EAPS domain.
new_name	Specifies a new name for the EAPS domain.

Default

N/A.

Usage Guidelines

None.

Example

The following command renames EAPS domain “eaps-1” to “eaps-5”:

```
config eaps eaps-1 name eaps-5
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the “*i*” series platforms.

create eaps

```
create eaps <name>
```

Description

Creates an EAPS domain with the specified name.

Syntax Description

name	Specifies the name of an EAPS domain to be created. May be up to 32 characters in length.

Default

N/A.

Usage Guidelines

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.

Example

The following command creates EAPS domain eaps_1 on an "*i*" series switch:

```
create eaps eaps-1
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

delete eaps

```
delete eaps <name>
```

Description

Deletes the EAPS domain with the specified name.

Syntax Description

name	Specifies the name of an EAPS domain to be deleted.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes EAPS domain eaps_1:

```
delete eaps eaps-1
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

disable eaps

```
disable eaps {<name>}
```

Description

Disables the EAPS function for a named domain or for an entire switch.

Syntax Description

name	Specifies the name of an EAPS domain.

Default

Disabled for the entire switch.

Usage Guidelines

None.

Example

The following command disables the EAPS function for entire switch:

```
disable eaps
```

The following command disables the EAPS function for the domain “eaps-1”:

```
disable eaps eaps-1
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the “*i*” series platforms.

enable eaps

```
enable eaps {<name>}
```

Description

Enables the EAPS function for a named domain or for an entire switch.

Syntax Description

name	Specifies the name of an EAPS domain.

Default

Disabled.

Default command enables for the entire switch.

Usage Guidelines

EDP must be enabled on the switch.

Example

The following command disables the EAPS function for entire switch:

```
enable eaps
```

The following command disables the EAPS function for the domain “eaps-1”:

```
enable eaps eaps-1
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the “*i*” series platforms.

show eaps

```
show eaps {<name>} {detail | summary}
```

Description

Displays EAPS status information.

Syntax Description

name	Specifies the name of an EAPS domain.
detail	Specifies all available detail for each domain.
summary	Specifies a summary of each domain.

Default

N / A.

Usage Guidelines

If you enter the `show eaps` command without a keyword, the command displays more information than with the `summary` keyword, but less than with the `detail` keyword.

Use the optional domain name parameter to display status information for a specific EAPS domain.

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 fields displayed are as follows:

EAPS Enabled:	Current state of EAPS on this switch: • Yes—EAPS is enabled on the switch. • no—EAPS is not enabled.
Number of EAPS instances:	Number of EAPS domains created. The maximum number of EAPS domains per switch is 64.
EAPSD-Bridge links:	The total number of EAPS bridge links in the system. The maximum count is 4096. Each time a VLAN is added to EAPS, this count increments by 1.
Name:	The configured name for this EAPS domain.
(instance=)	The instance number is created internally by the system.

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.
Preforwarding Timer interval:	The configured value of the timer. This value is set internally by the EAPS software. Displayed only for transit nodes.
Last update:	Displayed only for transit nodes; indicates the last time the transit node received a hello packet from the master node (identified by its MAC address).
EAPS Domain has … Controller VLANs:	Lists the assigned name and ID of the control VLAN.
EAPS Domain has … Protected VLANs:	Lists the assigned names and VLAN IDs of all the protected VLANs configured on this EAPS domain. Displayed only when the detail keyword is used.
Number of Protected VLANs:	The count of protected VLANs configured on this EAPS domain.

Example

The following command displays detailed EAPS information for domain “eaps2”:

```
show eaps eaps2 detail
```

The results for domain “eaps2” on a master node are shown as follows:

```
  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
```

The following command displays detailed EAPS information:

```
show eaps detail
```

The results for a transit node are shown as follows:

```
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 command displays EAPS information:

```
show eaps eaps2
```

The results for a transit node are shown as follows:

```
   Name: "eaps2" (instance=1)
   State: Link-Down       [Running: Yes]
   Enabled: Yes    Mode: Transit
   Primary port: 3           Port status: Down      Tag status: Tagged
   Secondary port: 2         Port status: Up        Tag status: Tagged
```

```
   Hello Timer interval: 1 sec     Fail Timer interval: 3 sec
   Preforwarding Timer interval: 6 sec
   Last update: From Master Id 00:01:30:B5:72:F0, at Fri Jun 7 23:13:09 2002
   EAPS Domain has following Controller Vlan:
      Vlan Name            VID             QosProfile
      "cv2"                4002            QP8
   Number of Protected Vlans: 2
```

The following command displays summary EAPS information:

```
show eaps summary
```

The results for this command are as follows:

```
EAPS Enabled: Yes
Number of EAPS instances: 2
EAPSD-Bridge links: 6

Domain        State        Mode     Enabled  Control-VLAN(VID)   #Protect-VLANs
------------  -----------  -------  -------  ------------------  --------------
eaps2         Link-Down    Transit  Yes      cv2         (4002)    2
eaps1         Link-Down    Transit  Yes      cv1         (4001)    2
```

History

This command was first available in ExtremeWare 6.2.

The summary option was added in ExtremeWare 6.2.2.

Platform Availability

This command is available on the "*i*" series platforms.

unconfig eaps port

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

Description

Sets the specified port's internal configuration state to INVALID.

Syntax Description

name	Specifies the name of an EAPS domain.
primary	Specifies that the primary port should be unconfigured.
secondary	Specifies that the secondary port should be unconfigured.

Default

N / A.

Usage Guidelines

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 detail` command to display the status information about the port.

Example

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

```
unconfig eaps eaps_1 primary port
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

12 Commands for Status Monitoring and Statistics

This chapter describes:

- Commands for configuring and managing the syslog
- Commands for enabling and disabling the syslog
- Commands for enabling and disabling NetFlow flow statistics collection
- Commands for configuring flow-collection port and filtering options
- Commands for configuring the flow-collector devices to which NetFlow datagrams are exported

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

  — 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 trouble shooting procedures.

  By default, log entries that are assigned a critical or warning level are considered static entries and remain in the log after a switch reboot.

- Subsystem

  The subsystem refers to the specific functional area to which the error refers.

  — 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.

The switch maintains 1,000 messages in its internal log. You can display a snapshot of the log at any time. 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. In addition to maintaining an internal log, the switch supports remote logging by way of the UNIX syslog host facility.

NetFlow Statistics

NetFlow flow statistics provides a way for a switch to capture and export traffic classification or precedence information as data traverses, or flows, across portions of a network. A network flow is defined as a unidirectional sequence of packets between a particular source device and destination device that share the same protocol and transport-layer information. Flows are defined by the combination of their source IP address, destination IP address, source port, destination port, and protocol type.

NetFlow records are unidirectional in nature, which means that two flow records are maintained for a typical TCP connection: one record for flow in the ingress direction; a second for the flow in the egress direction. Records are maintained only for TCP and UDP flows. Flow records are grouped together into UDP datagrams for export to a flow-collector device. A NetFlow Version 1 export datagram can contain up to 25 flow records.

The IP addresses (or hostnames) and UDP port numbers of the available flow collectors can be configured on a per-switch basis. The ExtremeWare NetFlow implementation also enables a single port to distribute statistics across multiple groups of flow-collector devices. The NetFlow distribution feature is enabled by configuring export distribution groups that contain the addresses of multiple flow-collector devices. The feature uses a distribution algorithm that ensures all of the records for a given flow are exported to the same collector. The algorithm also ensures that the flow records of the ingress direction of a TCP or UDP connection are exported to the same collector. For Ethernet applications, only ingress traffic is monitored on Ethernet ports.

By default, each Ethernet port configured for flow switching maintains statistics for all the flows traversing the link in the ingress direction. Generalized filtering options exist that enable you to configure an Ethernet port to maintain statistics selectively for only those flows that match a specified filter. Up to eight filters are supported for each Ethernet port, with a total of 128 filters possible per each I/O module.

NOTE

Some of the NetFlow commands are implemented differently in the version of ExtremeWare that supports the PoS module, than in ExtremeWare 6.2 or later. Commands or options unique to the PoS module are indicated in the comments, or are documented separately in Chapter 22.

clear counters

```
clear counters
```

Description

Clears all switch statistics and port counters, including port packet statistics, bridging statistics, and IP statistics.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

You should view the switch statistics and port counters before you clear them. Use the `show switch` command to view switch statistics. Use the `show port` command to view port statistics.

Viewing and maintaining 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. By clearing the counters, you can see fresh statistics for the time period you are monitoring.

Example

The following command clears all switch statistics and port counters:

```
clear counters
```

The following command clears (zeroes) all of the MPLS-related statistics:

```
clear counters
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in an ExtremeWare IP Technology Services Release based on v6.1.5b20 to support PoS modules.

This command was modified in an ExtremeWare IP Technology Services Release based on v6.1.8b12 to support MPLS modules.

Platform Availability

This command is available on all platforms.

clear log

```
clear log {diag-status | static}
```

Description

Clears the log database.

Syntax Description

diag-status	Diagnostic status level
static	If static is specified, the critical log messages are also cleared.

Default

N/A.

Usage Guidelines

The switch log tracks all configuration and fault information pertaining to the device.

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 `clear log static` command.

Example

The following command clears all log messages, including critical and warning log messages, from the database:

```
clear log static
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config flowstats export add port

```
config flowstats export <group#> add [<ipaddress> | <hostname>] port
<udp_port>
```

Description

Adds a flow-collector device to an export group to which NetFlow datagrams are exported.

Syntax Description

group#	Specifies the export group to which the specified flow-collector device should be added. The group number is an integer in the range of 1-32.
ipaddress	Specifies the IP address of a flow-collector destination.
hostname	Specifies the host name of a flow-collector destination.
udp_port	Specifies a UDP port for the destination flow-collector.

Default

N/A.

Usage Guidelines

You can configure up to 32 export distribution groups. Each group may contain up to eight flow-collection devices. At least one flow-collector destination must be configured for NetFlow datagrams to be exported to a group.

When multiple flow-collectors are configured as members of the same group, the exported NetFlow datagrams are distributed across the available destinations. This NetFlow-distribution feature enables a scalable collection architecture that is able to accommodate high volumes of exported data. The distribution algorithm ensures that all the records for a given flow are exported to the same collector. The algorithm also ensures that flow records for both the ingress and egress directions of a TCP or UDP connection are exported to the same collector (when both flows traverse the same link and both filters are configured to export to the same group).

Issuing this command also enables the collection of NetFlow statistics.

See Chapter 22 for information on a similar command for the PoS module (BlackDiamond 6800 only).

Example

The following command adds the flow-collector device with IP address 10.205.30.15 using UDP port 2025 to export group 5 for this switch:

```
config flowstats export 5 add 10.205.30.15 port 2025
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config flowstats export delete port

```
config flowstats export <group#> delete [<ipaddress> | <hostname>] port
<udp_port>
```

Description

Removes a flow-collector device from an export group to which NetFlow datagrams are exported.

Syntax Description

group#	Specifies the export group to which the specified flow-collector device belongs. The group number is an integer in the range of 1-32.
ipaddress	Specifies the IP address of the flow-collector destination.
hostname	Specifies the host name of the flow-collector destination.
udp_port	Specifies a UDP port of the destination flow-collector.

Default

N / A.

Usage Guidelines

See Chapter 22 for information on a similar command for the PoS module (BlackDiamond 6800 only).

Example

The following command removes the flow-collector device with IP address 10.205.30.15 using UDP port 2025 from export group 5 on this switch:

```
config flowstats export 5 delete 10.205.30.15 port 2025
```

History

This command first available in ExtremeWare 6.2 for "*i*" series platforms.

Platform Availability

This command is available on the "*i*" series platforms.

config flowstats filter-ingress ports export

```
config flowstats filter-ingress <filter#> ports <portlist> export <group#>
{aggregation} [<filterspec> | match-all-flows]
```

Description

Configures a flow record filter for the specified ports.

Syntax Description

filter#	Specifies the filter being defined. Value is an integer in the range of 1-8.
portlist	Specifies a list of ports to whose flows this filter should be applied. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
group#	Specifies the export group to which any matching flows should be exported.
aggregation	Specifies that all flows that match the filter should be maintained as a single set of statistics. This keyword is required if the filter is defined for a range of source or destination IP addresses and ports. Otherwise, it is optional.

filterspec	Specifies a set of five parameters (four are value/mask pairs) that define the criteria by which a flow is evaluated to determine if it should be exported. The parameters are:
	`destination [<ipaddress/ipaddress_mask> | any]`
	`ip-port [<portlist>/port_mask> | any]`
	`source [<ipaddress/ipaddress_mask> | any]`
	`ip-port [<portlist>/port_mask> | any]`
	`[ip | tcp | upd]`
	All five specifications must be included in the order specified.
	The range for port/port_mask is calculated using the following formula:
	(minport = port, maxport = 2^(32-port_mask)-1).

match-all-flows	Specifies that the filter should match any flow.

Default

N / A.

Usage Guidelines

Configuring a filter specification enables that filter for the specified ports. To specify all ports, you can use specify them as the range of all ports (such as 1-32 or 7:1-7:4) or in the form <slot>:* on a modular switch.

Each Ethernet port supports eight filters for ingress flows.

Conceptually, the filters work by ANDing the contents of each of the five components of a forwarded flow with the associated masks from the first defined filter (filter #1). Statistics are maintained if the results of the AND operations match the configured filter values for all fields of the sequence. If there is no match, then the operation is repeated for filter #2, and so on. If there is no match for any of the filters, then statistics are not maintained for the flow. Filters for any or all of the sequence components can be configured with a single command.

Example

The following command example configures filter 2 to collect aggregate statistics for all traffic flowing through ports 1-8 from the 192.170.0.0/16 subnet to the 192.171.132.0/24 subnet:

```
config flowstats filter-ingress 2 ports 1-8 export 1 aggregation destination
192.171.132.0/24 ip-port 0/0 source 192.170.0.0/16 ip-port 0/0 ip
```

The following command configures filter 3 to collect statistics on any flows for ports 4-32 that did not match the filters defined in filters 1 and 2:

```
config flowstats filter-ingress 3 ports 4-32 export 1 aggregation match-all-flows
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config flowstats source

```
config flowstats source <ipaddress>
```

Description

Configures the IP address that is to be used as the source IP address for NetFlow datagrams to be exported.

Syntax Description

ipaddress	Specifies the IP address of a VLAN to be used as the source address for the Net FL ow datagrams.

Default

Uses the IP address of the VLAN that has the default route to the flow-collector device.

Usage Guidelines

The IP address must have a route to the flow-collector device.

Example

The following command specifies that IP address 198.168.100.1 is the source:

```
config flowstats source 198.168.100.1
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config flowstats timeout ports

```
config flowstats timeout <minutes> ports [<portlist> | any]
```

Description

Configures the timeout value for flow records on the specified ports.

Syntax Description

minutes	Specifies the number of minutes to use in deciding when to export flow records. The default is five minutes.
portlist	Specifies the ports to which the timeout applies. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8. any indicates that the timeout should be set for all ports on this switch. Note: The parameter any is not supported for the PoS module.

Default

Five minutes.

Usage Guidelines

The timeout is used to export flow records on an age basis. All flow records are examined at least once every 30 minutes. If the age of the flow record is greater than the configured timeout, the record is exported. If the flow is still active, a new flow record will be created when the next packet arrives.

For the PoS module, the `minutes` parameter is an integer in the range [1-1440].

Example

The following command configures a timeout value of 15 minutes for ports 1-8:

```
config flowstats timeout 15 ports 1-8
```

This means that flow records for these ports will be exported after they have aged 15 minutes.

History

This command was first available in ExtremeWare 6.1.5b20 for the PoS module only

This command was first available in ExtremeWare 6.2 for all "*i*" series platforms.

Platform Availability

This command is available on the "*i*" series platforms.

config log display

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

Description

Configures the real-time log display.

Syntax Description

severity	Specifies a message severity. Severities include emergency, alert, critical, error, warning, info, notice, and debug.

Default

If not specified, only critical, alert, and emergency severity messages are sent to the syslog host.

Usage Guidelines

You must enable the log display before you can configure the log display. Use the `enable log display` command to enable the log display. This allows you to configure the system to maintain a running real-time display of log messages on the console.

Options for displaying the real-time log display include:

- severity — Filters the log to display messages with the selected severity or higher (more critical). Severities include emergency, alert, critical, error, warning, info, notice, and debug.

Example

The following command configures the system log to maintain a running real-time display of log messages of alert priority:

```
config log display alert
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config sys-health-check alarm-level

```
config sys-health-check alarm-level [log | system-down | traps |
auto-recovery <number of tries> [online | offline]]
```

Description

Configures the system health checker.

Syntax Description

log	Posts a CRIT message to the log.
system-down	Posts a CRIT message to the log, sends a trap, and turns off the system.
traps	Posts a CRIT message to the log and sends a trap.
auto-recovery	Specifies the number of times that the health checker attempts to auto-recover. The range is from 0 through 255 times. Default is 3 times.
offline	Specifies that a faulty Summit switch or Alpine module is taken offline and kept offline if one of the following occurs: • More than eight defects are detected. • Three consecutive checksum errors were detected by the health checker, but no new defects were found by the memory scanning and mapping process. • After defects were detected and mapped out, the same checksum errors are again detected by the system health checker.
online	Specifies that a faulty module is kept online, regardless of how many errors are detected.

Default

The default alarm level is log.

Usage Guidelines

This command allows you to configure the switch's reaction to a failed health check.

The system health checker tests I/O modules, SMMi modules, and the backplane by forwarding packets every 4 seconds. Additional checking for the validity of these packets is completed by performing a checksum.

By isolating faults to a specific module, SMMi, or backplane connection, the system health checker prevents corrupted packets from being propagated to the CPU, upper layer modules, or the rest of your network. If you observe a failure, please contact Extreme Technical Support.

Depending on how you have configured the system health checker, failed system health checks may generate critical error messages in the syslog, and may also send a trap and/or shut down the system. The system health checker will continue to periodically forward test packets to failed components. If auto-recovery is configured, the system will attempt to automatically reset the faulty module and bring it back online.

The alarm-level and auto-recovery options are mutually exclusive; configuring an alarm-level disables auto-recovery, and configuring auto-recovery overrides the alarm-level setting.

In ExtremeWare versions prior to 6.2, you cannot use both mirroring and the system health checker at the same time. If you configure mirroring with the system health checker enabled, the health checker will indicate that it has been disabled by sending a message to the syslog. In ExtremeWare 6.2 or later, this restriction does not apply.

The `auto-recovery` option configures the number of times the system health checker attempts to automatically reset a faulty module and bring it online. If the system health checker fails more than the configured number of attempts, it sets the module to card-down.

In ExtremeWare 6.2.1 or later, when auto-recovery is configured, the occurrence of three consecutive checksum errors will cause the packet memory (PM) defect detection program to be run against the I/O module. Checksum errors may include internal and external MAC port parity errors, EDP checksum errors, and CPU packet or diagnostic packet checksum errors. If defects are detected, the card is taken offline, the memory defect information is recorded in the card EEPROM, the defective buffer is mapped out of further use, and the card is returned to operational state. A maximum of 8 defects can be stored in the EEPROM.

After the PM defect detection and mapping process has been run, a card is considered failed and is taken offline in the following circumstances:

- More than eight defects are detected.
- Three consecutive checksum errors were detected by the health checker, but no new PM defects were found by the PM defect detection process.
- After defects were detected and mapped out, the same checksum errors are again detected by the system health checker.

The auto-recovery repetition value is ignored in these cases. In any of these cases, please contact Extreme Technical Support.

To view the status of the system health checker, use the `show diag` command.

To enable the health checker, use the `enable sys-health-check` command.

To disable the health checker, use the `disable sys-health-check` command.

The alarm-level `system-down` option is especially useful in an ESRP configuration where the entire system is backed by an identical system. By powering down the faulty system, you ensure that erratic ESRP behavior in the faulty system does not affect ESRP performance and ensures full system failover to the redundant system.

If you are using ESRP in your configuration, any system health check failure will automatically reduce the ESRP priority of the system to the configured failover priority. This allows the healthy standby system to take over ESRP and become responsible for handling traffic.

If you specify the `online` option, the module is kept online, but the following error messages are recorded in the log:

```
<WARN:SYST> card_db.c 832: Although card 2 is back online, contact Tech. Supp. for
assistance.
<WARN:SYST> card_db.c 821: Card 2 has nonrecoverable packet memory defect.
```

I/O module faults are permanently recorded on the module's EEPROM. A module that has failed a system health check cannot be brought back online.

To view the failure messages, use the `show diag` command.

Example

The following command configures the system health checker to post a CRIT message to the log and send a trap:

```
config sys-health-check alarm-level traps
```

History

This command was first available in ExtremeWare 6.1.9.

The system health check functionality was modified in ExtremeWare 6.2.1 to support packet memory defect detection and mapping on selected I/O modules.

This command was modified in ExtremeWare 6.2.2 to support system health check and checksum error checking and to add the `online` and `offline` parameters.

Platform Availability

This command is available only on Alpine and Summit switches.

config sys-health-check auto-recovery

```
config sys-health-check auto-recovery <number of tries> [offline | online]
```

Description

Configures the system health checker.

Syntax Description

number of tries	Specifies the number of times that the health checker attempts to auto-recover a faulty module. The range is from 0 through 255 times. Default is 3 times.
offline	Specifies that a faulty module is taken offline and kept offline if one of the following occurs: • More than eight defects are detected. • Three consecutive checksum errors were detected by the health checker, but no new defects were found by the memory scanning and mapping process. • After defects were detected and mapped out, the same checksum errors are again detected by the system health checker.
online	Specifies that a faulty module is kept online, regardless of memory scanning or memory mapping errors.

Default

Log.

Usage Guidelines

This command allows you to configure the switch's reaction to a failed health check.

The system health checker tests I/O modules, MSM64i modules, and the backplane by forwarding packets every 4 seconds. Additional checking for the validity of these packets is completed by performing a checksum.

By isolating faults to a specific module, MSM64i, or backplane connection, the system health checker prevents corrupted packets from being propagated to the CPU, upper layer modules, or the rest of your network. If you observe a failure, please contact Extreme Technical Support.

Depending on how you have configured the system health checker, failed system health checks may generate critical error messages in the syslog, and may also send a trap and/or shut down the system. The system health checker will continue to periodically forward test packets to failed components. If auto-recovery is configured, the system will attempt to automatically reset the faulty module and bring it back online.

In ExtremeWare versions prior to 6.2, you cannot use both mirroring and the system health checker at the same time. If you configure mirroring with the system health checker enabled, the health checker will indicate that it has been disabled by sending a message to the syslog. In ExtremeWare 6.2 or later, this restriction does not apply.

The auto-recovery option configures the number of times the system health checker attempts to automatically reset a faulty module and bring it online. If the system health checker fails more than the configured number of attempts, it sets the module to card-down.

In ExtremeWare 6.2.1 or later, when auto-recovery is configured, the occurrence of three consecutive checksum errors will cause the packet memory (PM) defect detection program to be run against the I/O module. Checksum errors may include internal and external MAC port parity errors, EDP checksum errors, and CPU packet or diagnostic packet checksum errors. If defects are detected, the card is taken offline, the memory defect information is recorded in the card EEPROM, the defective buffer is mapped out of further use, and the card is returned to operational state. A maximum of 8 defects can be stored in the EEPROM.

After the PM defect detection and mapping process has been run, a card is considered failed and is taken offline in the following circumstances:

- More than eight defects are detected.
- Three consecutive checksum errors were detected by the health checker, but no new defects were found by the memory scanning and mapping process.
- After defects were detected and mapped out, the same checksum errors are again detected by the system health checker.

The auto-recovery repetition value is ignored in these cases. In any of these cases, please contact Extreme Technical Support.

Auto-recovery mode only affects an MSM64i if the system has no slave MSM64i. If the faulty module is the only MSM64i in the system, auto recovery automatically resets the MSM64i and brings it back online. Otherwise, auto-recovery has no effect on an MSM64i.

To view the status of the system health checker, use the `show diag` command.

To enable the health checker, use the `enable sys-health-check` command.

To disable the health checker, use the `disable sys-health-check` command.

The alarm-level `system-down` option is especially useful in an ESRP configuration where the entire system is backed by an identical system. By powering down the faulty system, you ensure that erratic ESRP behavior in the faulty system does not affect ESRP performance and ensures full system failover to the redundant system.

If you are using ESRP in your configuration, any system health check failure will automatically reduce the ESRP priority of the system to the configured failover priority. This allows the healthy standby system to take over ESRP and become responsible for handling traffic.

If you specify the `online` option, the module is kept online, but the following error messages are recorded in the log:

```
<WARN:SYST> card_db.c 832: Although card 2 is back online, contact Tech. Supp. for
assistance.
<WARN:SYST> card_db.c 821: Card 2 has nonrecoverable packet memory defect.
```

I/O module faults are permanently recorded on the module's EEPROM. A module that has failed a system health check cannot be brought back online.

If the faulty module is a master MSM64i, the slave MSM64i automatically becomes the master and sets the faulty MSM64i to `card-down`. The new master MSM64i re-initializes and brings up all the I/O modules.

If the faulty module is a master MSM64i and there is no slave MSM64i, the system continues operation in a "limited commands" mode. In the "limited commands" mode, the I/O slots are not initialized, and only commands that do no affect the switch hardware configuration are allowed.

If the faulty module is a slave MSM64i, the fault is recorded in the slave's MSM64i's NVRAM and the slave MSM64i is taken offline.

To view the failure messages, use the `show diag` command.

To clear the MSM64i failure messages posted to the log, use the `clear log diag-status` command. This command will clear the error messages from the MSM64i NVRAM. If the MSM64i failed a system health check, this command restores the MSM64i to full functionality. This command should only be used for additional testing purposes and reproduction efforts of the original fault.

Example

The following command configures the system health checker to try ten times to bring a faulty MSM64i back online:

```
config sys-health-check auto-recovery 10
```

History

This command was first available in ExtremeWare 6.1.9.

The system health check functionality was modified in ExtremeWare 6.2.1 to support packet memory defect detection and mapping on selected I/O modules.

This command was modified in ExtremeWare 6.2.2 to support system health check and checksum error checking on the BlackDiamond 6804 and to add the `online` and `offline` parameters.

Platform Availability

This command is available only on BlackDiamond chassis.

The packet-memory defect detection and mapping feature is supported only on selected I/O modules. See the Release Note for your version of ExtremeWare for information on the supported modules.

config syslog

```
config syslog {add} [<hostname> | <ip_address>] <facility> {<severity>}
```

Description

Configures the syslog host address, and filters messages to be sent to the syslog host.

Syntax Description

hostname	Species the hostname of the syslog host.
ip_address	Specifies an IP address of the syslog host.
facility	Specifies a syslog facility level.
severity	Specifies a message severity. Severities include emergency, alert, critical, error, warning, info, notice, and debug.

Default

If a severity level is not specified, all messages are sent to the syslog host.

Usage Guidelines

Options for configuring the syslog include:

- hostname/ipaddress — The name or IP address of the syslog host.
- facility — The syslog facility level for local use (local0– local7).
- severity — Filters the log to display messages with the selected severity or higher (more critical). Severities include emergency, alert, critical, error, warning, info, notice, and debug.

The switch syslog overwrites existing log information in a wrap-around buffer, which may cause you to lose valuable information once the buffer becomes full. The syslog host does not overwrite log information.

The enable syslog command must be issued in order for messages to be sent to the syslog server(s). Syslog is disabled by default. A total of four syslog servers can be configured at one time.

For version 4.0 and higher:

- The syslog facility level is defined as local0 – local7. The facility level is used to group syslog data.

Example

The following command configures remote logging with an emergency priority:

```
config syslog 123.45.67.78 level1 emergency
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config syslog delete

```
config syslog delete [<hostname> | <ip_address>] <facility> {<severity>}
```

Description

Deletes a syslog host address.

Syntax Description

hostname	Species the hostname of the syslog host.
ip_address	Specifies an IP address of the syslog host.
facility	Specifies a syslog facility level.
severity	Specifies a message severity. Severities include emergency, alert, critical, error, warning, info, notice, and debug.

Default

N/A.

Usage Guidelines

This command is used to delete a syslog host.

Example

The following command deletes the syslog host with an IP address of 10.0.0.1:

```
config syslog delete 10.0.0.1 local1 alert
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on all platforms.

disable cli-config-logging

```
disable cli-config-logging
```

Description

Disables the logging of CLI configuration commands to the switch Syslog.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

The `disable cli-config-logging` command discontinues the recording of all switch configuration changes and their sources that are made using the CLI via Telnet or the local console. After you disable configuration logging, no further changes are logged to the system log.

To view the status of configuration logging on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for configuration logging.

Example

The following command disables the logging of CLI configuration command to the Syslog:

```
disable cli-config-logging
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on all platforms.

disable flowstats

```
disable flowstats
```

Description

Disables the flow statistics feature on the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

When this feature is disabled, no flow records are exported.

Example

The following command disables the NetFlow statistics feature on this switch:

```
disable flowstats
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

disable flowstats filter ports

```
disable flowstats filter <filter#> ports <portlist> [ingress | egress]
```

Description

Disables a specified flow record filter for the specified ports.

Syntax Description

filter#	Specifies the flow record filter that should be disabled.
portlist	Specifies a list of ports or slots and ports for which the filter should be disabled. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
ingress	Use this keyword to specify that the filter being enabled or disabled is one of the eight filters to be applied to inbound flows. Supported on the PoS module only.
egress	Use this keyword to specify that the filter being enabled or disabled is one of the eight filters to be applied to outbound flows. Supported on the PoS module only.

Default

For the PoS module, filter #1 is enabled on all SONET ports, and the remaining filters are disabled.

For other switches or modules, filters are enabled by default when they are configured.

Usage Guidelines

The `filter#` parameter is an integer in the range [1-8].

For each SONET port on a PoS module, sixteen filters are supported—eight filters for ingress flows and another eight filters for egress flows. The `filter#` parameter and either the `ingress` or `egress` keyword are used to identify the particular filter that is being disabled.

One of either the `ingress` or `egress` keywords are required for SONET ports.

Example

The following command disables filter 3 for ports 1-8 on an "*i*" series switch:

```
disable flowstats filter 3 ports 1-8
```

The following command example disables ingress filter #2 on port 1 of the PoS module installed in slot 8 of the BlackDiamond switch:

```
disable flowstats filter 2 ports 8:1 ingress
```

History

This command was first available in ExtremeWare 6.1.5b20 for the PoS module only.

This command was first available in ExtremeWare 6.2 for all "*i*" series platforms.

Platform Availability

This command is available on the "*i*" series platforms.

disable flowstats ping-check

```
disable flowstats ping-check <group#>
```

Description

Disables the flow statistics ping-check function for a specified group of collector devices.

Syntax Description

group#	Specifies the export group for which the ping-check function should be disabled.

Default

Enabled.

Usage Guidelines

On the PoS module, if you do not include a group number, ping-check is disabled for all export groups. The group number is not optional for other Extreme "*i*" series devices.

Example

The following command disables the ping-check function for export group 3.

```
disable flowstats ping-check 3
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

disable flowstats ports

```
disable flowstats ports <portlist>
```

Description

Disables the flow statistics function on the specified ports.

Syntax Description

portlist	Specifies a list of ports or slots and ports for which the flowstats function should be disabled. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

N/A.

Usage Guidelines

On the PoS module, flow statistics are only collected on SONET ports that are configured to use the IP control protocol, IPCP, (in other words, flow statistics are not collected on ports that are configured to use the bridging control protocol, BCP). Also, there are no configuration restrictions that prohibit enabling of the flow statistics function on ports that are not configured to use IPCP; statistics are not collected on those ports.

Example

The following command disables NetFlow statistics for ports 1-8 on this switch:

```
disable flowstats ports 1-8
```

History

This command was first available in ExtremeWare 6.1.5b20 for the PoS module

This command was made available in ExtremeWare 6.2 for all “i” series platforms.

Platform Availability

This command is available on the “*i*” series platforms.

disable log display

```
disable log display
```

Description

Disables the log display.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

If the log display is disabled, log information is no longer written to the serial console.

Example

The following command disables the log display:

```
disable log display
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

disable rmon

```
disable rmon
```

Description

Disables the collection of RMON statistics on the switch.

Syntax Description

This command has no arguments or variables.

Default

By default, RMON is disabled. However, even in the disabled state, the switch responds to RMON queries and sets for alarms and events.

Usage Guidelines

The switch supports four out of nine groups of Ethernet RMON statistics. In a disabled state, the switch continues to respond to the following two groups:

- 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 may be autocalibrated or set manually.
- 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.

To view the status of RMON polling on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for RMON polling.

Example

The following command disables the collection of RMON statistics on the switch:

```
disable rmon
```

History

This command was first available in ExtremeWare 4.1.

Platform Availability

This command is available on all platforms.

disable sys-backplane-diag

```
disable sys-backplane-diag
```

Description

Disables system run time backplane diagnostics.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

If the system backplane diagnostics is disabled, it does not check for backplane faults.

Example

The following command disables system run time backplane diagnostics:

```
disable sys-backplane-diag
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on BlackDiamond switches only.

disable sys-health-check

```
disable sys-health-check
```

Description

Disables the BlackDiamond system health checker.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

If the system health checker is disabled, it does not test I/O modules, MSM64i modules, and the backplane for system faults.

Example

The following command disables the BlackDiamond system health checker:

```
disable sys-health-check
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on BlackDiamond switches only.

disable syslog

```
disable syslog
```

Description

Disables logging to a remote syslog host.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Disables logging to a remote syslog host, not to the switch syslog.

Example

The following command disables logging to a remote syslog host:

```
disable syslog
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

enable cli-config-logging

```
enable cli-config-logging
```

Description

Enables the logging of CLI configuration commands to the Syslog for auditing purposes.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

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 changes and the source IP address of the client (if Telnet was used). Configuration logging applies only to commands that result in a configuration change.

Every command is displayed in the log window. This allows you to view every command executed on the switch.

To view the status of configuration logging on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for configuration logging.

Example

The following command enables the logging of CLI configuration commands to the Syslog:

```
enable cli-config-logging
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on all platforms.

enable flowstats

```
enable flowstats
```

Description

Enables the flow statistics feature on the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables NetFlow statistics feature on this switch:

```
enable flowstats
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

enable flowstats filter ports

```
enable flowstats filter <filter#> ports <portlist> {ingress | egress}
```

Description

Enables a specified flow record filter for the specified ports.

Syntax Description

filter#	Specifies the flow record filter that should be enabled.
portlist	Specifies the ports or slots and ports for which the filter should be enabled. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
ingress	Use this keyword to specify that the filter being enabled or disabled is one of the eight filters to be applied to inbound flows on the SONET port(s). Supported on the PoS module only.
egress	Use this keyword to specify that the filter being enabled or disabled is one of the eight filters to be applied to outbound flows on the SONET port(s). Supported on the PoS module only.

Default

For the PoS module, filter #1 is enabled on all SONET ports, and the remaining filters are disabled.

For other switches or modules, filters are enabled by default when they are configured.

Usage Guidelines

The `filter#` parameter is an integer in the range [1-8]. A filter must be enabled to match a flow. For "i" series devices other than the PoS module, these apply to outbound flows only.

For each SONET port on a PoS module, sixteen filters are supported—eight filters for ingress flows and another eight filters for egress flows. The `filter#` parameter and either the `ingress` or `egress` keyword are used to identify the particular filter that is being disabled.

One of either the `ingress` or `egress` keywords are required for SONET ports.

Example

The following command enables filter 3 for ports 1-8 on the switch:

```
enable flowstats filter 3 ports 1-8
```

History

This command was first available in ExtremeWare 6.1.5b20 for the PoS module only.

This command was first available in ExtremeWare 6.2 for all "*i*" series platforms.

Platform Availability

This command is available on the "*i*" series platforms.

enable flowstats ping-check

```
enable flowstats ping-check <group#>
```

Description

Enables the flow statistics ping-check function for a specified group of collector devices.

Syntax Description

group#	Specifies the export group for which the ping-check function should be enabled.

Default

Enabled.

Usage Guidelines

If a flow-collector device is repeatedly unresponsive to ping requests, it is temporarily removed from the distribution list for any export groups of which it is a member. The device will be returned to the distribution list automatically when subsequent ping-checks are successful.

On the PoS module, if you do not include a group number, ping-check is enabled for all export groups.

Example

The following command enables the ping-check function for export group 3.

```
enable flowstats ping-check 3
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

enable flowstats ports

```
enable flowstats ports <portlist>
```

Description

Enables the flow statistics function on the specified ports.

Syntax Description

portlist	Specifies a list of ports or slots and ports for which the flowstats function should be enabled. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

Disabled.

Usage Guidelines

On the PoS module, flow statistics are only collected on SONET ports that are configured to use the IP control protocol, IPCP, (in other words, flow statistics are not collected on ports that are configured to use the bridging control protocol, BCP). Also, there are no configuration restrictions that prohibit enabling of the flow statistics function on ports that are not configured to use IPCP; statistics are not collected on those ports.

Example

The following command enables NetFlow statistics for ports 1-8 on this switch:

```
enable flowstats ports 1-8
```

History

This command was first available in ExtremeWare 6.1.5b20 for the PoS module

This command was made available in ExtremeWare 6.2 for all “i” series platforms.

Platform Availability

This command is available on the “*i*” series platforms.

enable log display

```
enable log display
```

Description

Configures the system to maintain a running real-time display of log messages on the console.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

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.

You configure the commands displayed in the log using the `config log display` command.

Example

The following command enables a real-time display of log messages:

```
enable log display
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

enable rmon

```
enable rmon
```

Description

Enables the collection of RMON statistics on the switch.

Syntax Description

This command has no arguments or variables.

Default

By default, RMON is disabled. However, even in the disabled state, the switch responds to RMON queries and sets for alarms and events. By enabling RMON, the switch begins the processes necessary for collecting switch statistics.

Usage Guidelines

The switch supports four out of nine groups of Ethernet RMON statistics. In an enabled state, the switch responds to the following four groups:

- Statistics—The RMON Ethernet Statistics group provides traffic and error statistics showing packets, bytes, broadcasts, multicasts, and errors on a LAN segment or VLAN.
- 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.
- 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 may be autocalibrated or set manually.
- 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.

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.

RMON requires one probe per LAN segment, and standalone RMON probes have traditionally 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 view the status of RMON polling on the switch, use the `show management` command. The `show management` command displays information about the switch including the enable/disable state for RMON polling.

Example

The following command enables the collection of RMON statistics on the switch:

```
enable rmon
```

History

This command was first available in ExtremeWare 4.1.

Platform Availability

This command is available on all platforms.

enable sys-backplane-diag

```
enable sys-backplane-diag
```

Description

Enables system run time backplane diagnostics, which is done by periodically sending diagnostic packets between MSM and I/O modules while the system is running.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

If the system health check detects a backplane fault in a module, the system automatically reconfigures the backplane link map to redistribute traffic over the remaining backplane links. If all backplane links to a module fail, the module is considered down.

To view the status of the links between the modules and each MSM64i, use the `show internal-port-stats slot <slot_num>` command. Where the slot_num is the number of the slot that you want to view the status of the links between the modules.

Example

The following command enables system run time backplane diagnostics on the BlackDiamond switch:

```
enable sys-backplane-diag
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on BlackDiamond switches only.

enable sys-health-check

```
enable sys-health-check
```

Description

Enables the BlackDiamond system health checker.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

The system health checker tests I/O modules, MSM64i modules, and the backplane by forwarding packets every 4 seconds. Additional checking for the validity of these packets is completed by performing a checksum.

By isolating faults to a specific module, MSM64i, or backplane connection, the system health checker prevents corrupted packets from being propagated to the CPU, upper layer modules, or the rest of your network. If you observe a failure, please contact Extreme Technical Support.

Depending on how you have configured the system health checker, failed system health checks may generate critical error messages in the syslog, and may also send a trap and/or shut down the system. The system health checker will continue to periodically forward test packets to failed components. If auto-recovery is configured, the system will attempt to automatically reset the faulty module and bring it back online.

In ExtremeWare versions prior to 6.2, you cannot use both mirroring and the system health checker at the same time. If you configure mirroring with the system health checker enabled, the health checker will indicate that it has been disabled by sending a message to the syslog. In version 6.2 or later, this restriction does not apply.

To configure the health checker, use the following command:

```
config sys-health-check [alarm-level [card-down | default | log | system-down | traps]
| auto-recovery <number of tries>]
```

The alarm-level and auto-recovery options are mutually exclusive; configuring an alarm-level disables auto-recovery, and configuring auto-recovery overrides the alarm-level setting.

Example

The following command enables the BlackDiamond system health checker:

```
enable sys-health-check
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on BlackDiamond switches only.

enable syslog

```
enable syslog
```

Description

Enables logging to a remote syslog host.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

In order to enable remote logging, you must do the following:

- Configure the syslog host to accept and log messages.
- Enable remote logging by using the `enable syslog` command.
- Configure remote logging by using the `config syslog` command.

When you use the `enable syslog` command, the exporting process of the syslog begins.

Example

The following command enables logging to a remote syslog host:

```
enable syslog
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

show flowstats

```
show flowstats {detail}
```

Description

Displays status information for the flow statistics function.

Syntax Description

detail	Specifies that detailed NetFlow configuration information should be displayed.

Default

Displays summary statistics information for all ports.

Usage Guidelines

The command with no arguments displays flowstats configuration information for all ports. The information is displayed in a format similar to the flowstats command syntax. For the statistics that apply to individual ports, the port number is presented without a "port" keyword. For example, in the `NetFlow Enable/Disable per port` and `NetFlow TimeOut Config` sections of the example below, the port number immediately follows the `flowstats` keyword. The following information is displayed:

- Whether the flowstats feature is enabled or disabled
- Whether flowstats is enabled or disabled for individual ports
- The configuration of flow-collector devices (`NetFlow Server Config`)
- NetFlow Timeout configurations
- Whether NetFlow Filters are enable or disabled
- NetFlow filter specifications
- NetFlow ping-check configuration

When the `detail` keyword is included, the NetFlow Servers Config section is replaced by detailed configuration information that includes counts of the number of times each flow-collector device has been removed from the distribution list due to ping-check failures.

For each export group, the following information is displayed:

- Whether ping-check is enabled
- The source IP address
- An entry for each flow-collector device in the export group, displaying the following:
  - The IP address of the device
  - The UPD port number for the device
  - Whether the device is up or down (based on the ping-check response)
  - The number of times the device has been unreachable based on the ping-check response

Example

The `show flowstats` command with no options, for a switch with NetFlow statistics enabled on ports 1, 40, and 43, displays output similar to the following:

```
* Summit48i:1 # show flowstats
#
enable flowstats
#
# NetFlow Enable/Disable per port
enable flowstats 1
enable flowstats 40
enable flowstats 43
#
# NetFlow Servers Config
config flowstats export 1 add 10.0.6.40 udp_port 3333
config flowstats export 1 add 10.0.1.1 udp_port 9999
config flowstats export 3 add 10.0.6.70 udp_port 1212
config flowstats export 32 add 10.0.6.40 udp_port 6666
#
# NetFlow TimeOut Config
config flowstats 1 timeout 5
config flowstats 40 timeout 5
config flowstats 43 timeout 5
#
#
# NetFlow Filters Enable/Disable
enable flowstats 1 filter 1
enable flowstats 40 filter 8
enable flowstats 43 filter 3
#
# NetFlow Filter Specifications Config
config flowstats 1 filter-ingress 1 export 3 aggregation dest_ip 10.203.0.1/ffff
ffff source_ip 0.0.0.0/0 dest_port 0/0 source_port 0/0 protocol IP
config flowstats 40 filter-ingress 8  export 1 aggregation match-all-flows
config flowstats 43 filter-ingress 3 export 32 aggregation dest_ip 10.0.1.1/ffff
ffff source_ip 10.201.32.10/ffffffff dest_port 0/0 source_port 0/0 protocol TCP
#
# NetFlow Ping Check Config
```

The `show flowstats` command with the `detail` option, for a switch with NetFlow statistics enabled on ports 1, 40, and 43, displays output similar to the following:

```
* Summit48i:1 # show flowstats detail

enable flowstats
#

# NetFlow Enable/Disable per port
enable flowstats 1
enable flowstats 40
enable flowstats 43
#
# NetFlow Servers Config
Group: 1  ping-check: enable  Source ip_address: 0.0.0.0
--------------------------------------------------------------------------------
ip_address 10.0.6.40       udp_port  3333 status    up    0
```

```
ip_address 10.0.1.1        udp_port  9999 status    up    0
--------------------------------------------------------------------------------
Group: 3  ping-check: enable  Source ip_address: 0.0.0.0
--------------------------------------------------------------------------------
ip_address 10.0.6.70       udp_port  1212 status  down    0
--------------------------------------------------------------------------------
Group: 32  ping-check: enable  Source ip_address: 0.0.0.0
--------------------------------------------------------------------------------
ip_address 10.0.6.40       udp_port  6666 status    up    0
--------------------------------------------------------------------------------
```

The remaining output is the same as the previous example, starting with the `NetFlow Timeout Config` section.

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

show flowstats group

```
show flowstats group <group#>
```

Description

Displays configuration information an export group.

Syntax Description

group#	Specifies an export group for which configuration information should be displayed.

Default

N / A.

Usage Guidelines

The information displayed by this command is displayed in a format similar to the `config flowstats export` command. The following information is displayed:

- Whether the flowstats feature is enabled or disabled
- The configuration of flow-collector devices for the export group (`NetFlow Server Config`)
- NetFlow ping-check configuration

Example

The following command displays statistics for export group 3:

```
* Summit48i:4 # show flowstats group 3

enable flowstats
#
# NetFlow Servers Config
config flowstats export 3 add 10.0.6.70 udp_port 1212
#
#
# NetFlow Ping Check Config
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

show flowstats ports

```
show flowstats ports <portlist>
```

Description

Displays status information for the flow statistics function.

Syntax Description

portlist	Specifies a list of ports or slots and ports for which flow statistics should be displayed. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

N / A.

Usage Guidelines

This command displays flowstats configuration information for an individual port. The information is displayed in a format similar to the flowstats command syntax. The following information is displayed:

- Whether the flowstats feature is enabled or disabled
- Whether flowstats is enabled or disabled for the specified port
- NetFlow Timeout configuration for the port
- Whether NetFlow Filters are enable or disabled for the port
- NetFlow filter specifications for the port

Example

The following command displays statistics for port 40:

```
* Summit48i:4 # show flowstats ports 40
#
enable flowstats
#
# NetFlow Enable/Disable per port
enable flowstats 40
#
# NetFlow TimeOut Config
config flowstats 40 timeout 5
#
# NetFlow Filters Enable/Disable
enable flowstats 40 filter 8
#
# NetFlow Filter Specifications Config
config flowstats 40 filter-ingress 8  export 1 aggregation match-all-flows
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

show log

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

Description

Displays the current snapshot of the log.

Syntax Description

chronological	Specifies to show log entries in ascending chronological order.
priority	Specifies a message priority. These can be one of the following: alert, critical, debug, emergency, error, info, notice, or warning.
time	Specifies that the log entries should be displayed in chronological order, based on the timestamp.

Default

If not specified, informational priority messages and higher are displayed.

Usage Guidelines

The switch maintains 1,000 message in its internal log. You use the `show log` command to display the log.

Options for displaying the log include:

- `chronological`—Filters the log to display messages in ascending chronological order.
- `priority`—Filters the log to display messages with the selected priority or higher (more critical). Priorities include alert, critical, debug, emergency, error, info, notice, and warning.
- `time`—Filters the log to display messages based on the timestamp.

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 14 describes the levels of importance that the system can assign to a fault.

Table 14: 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 15 describes the subsystems.

Table 15: 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.

Example

The following command displays messages with a critical priority:

```
show log critical
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 6.2.2 to include the `chronological` option.

Platform Availability

This command is available on all platforms.

show log config

```
show log config
```

Description

Displays the log configuration.

Syntax Description

This command has no arguments or variables.

Default

N/A

Usage Guidelines

The `show log config` command displays the log configuration including the syslog host IP address, the priority level of messages being logged locally and whether the console log is enabled or disabled, and the priority level of messages being sent to the syslog host and whether the syslog is enabled or disabled.

Example

The following command displays the log configuration:

```
show log config
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

show memory

```
show memory {detail}
```

Description

Displays the current system memory information.

Syntax Description

detail	Specifies task-specific memory usage.

Default

N / A.

Usage Guidelines

Your BlackDiamond or Summit switch must have 32MB of DRAM to support the features in ExtremeWare version 4.0 and above.

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. This way, statistics can help you get the best out of your network.

The `show memory` command displays the following information in a tabular format:

- System memory information including the total DRAM size of your system.
- Current memory (both free and allocated memory) used by the system and the users.
- Cumulative memory (both free and allocated memory) used by the users.
- Software packet memory statistics including the type of packet, the number of allocated and free packets, the number of packet failures, and data and other blocks.
- Memory utilization statistics including the total blocks of memory available and the memory being used on your system. You can review how your memory is being utilized For example you can view memory utilization for the system, management, ESRP, IP, and other system functions.

This information may be useful for your technical support representative if you experience a problem.

For version 2.0 and 4.0:

- The `detail` parameter is not available.

Depending on the software version running on your switch, additional or different memory information may be displayed.

Example

The following command displays current system memory information:

```
show memory
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

show ports rxerrors

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

Description

Displays real-time receive error statistics.

For PoS modules, displays the `rxerror` information for the PoS ports. Only a subset of the statistics displayed by this command are applicable to PoS ports. The fields that do not apply to PoS ports are displayed with values of all zeroes.

Syntax Description

portlist	Specifies a list of ports or slots and ports to which the parameters apply. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

N / A.

Usage Guidelines

If you do not specify a port number or range of ports, receive error statistics are displayed for all ports.

This status information may be useful for your technical support representative if you have a network problem.

The following port receive error information is collected by the switch:

- Port Number
- Link Status — The current status of the link. Options are:
  - Ready (R): The port is ready to accept a link.
  - Active (A): The link is present at this port.
  - Disabled (D): The link is disabled at this port.
  - Not Present (NP): The link is not present at this port.
- Receive Bad CRC Frames (RX CRC) — The total number of frames received by the port that were of the correct length, but contained a bad FCS value.
- Receive Oversize Frames (RX Over) — The total number of good frames received by the port greater than the supported maximum length of 1,522 bytes. For products that use the “i” chipset, ports with jumbo frames enabled do no increment this counter.
- Receive Undersize Frames (RX Under) — The total number of frames received by the port that were less than 64 bytes long.
- Receive Fragmented Frames (RX Frag) — The total number of frames received by the port were of incorrect length and contained a bad FCS value.
- Receive Jabber Frames (RX Jabber) — 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.

For version 2.0 and 4.0

- Disabled and Not Present are not available as link status indicators.

Example

The following command displays receive error statistics for ports 1 through 3 on a stand-alone switch:

```
show ports 1-3 rxerrors
```

The following command displays receive error statistics for slot 1, ports 1 through 3 on a modular switch:

```
show ports 1:1-1:3 rxerrors
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 6.1 to support PoS modules.

This command was modified in ExtremeWare 6.0 to support the Disabled and Not Present link status indicators.

This command was modified in ExtremeWare 4.0 to support modular switches.

Platform Availability

This command is available on all platforms.

show ports stats

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

Description

Displays real-time port statistics.

Syntax Description

portlist	Specifies a list of ports or slots and ports to which the parameters apply. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

N/A.

Usage Guidelines

If you do not specify a port number or range of ports, statistics are displayed for all ports.

Jumbo frame statistics are displayed for *"i"* series switches only that are configured for jumbo frame support.

This status information may be useful for your technical support representative if you have a network problem.

The following port statistic information is collected by the switch:

- Port Number
- Link Status — The current status of the link. Options are:
  - Ready (R): The port is ready to accept a link.
  - Active (A): The link is present at this port.
  - Disabled (D): The link is disabled at this port.
  - Not Present (NP): The link is not present at this port.
- 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.

For version 2.0 and 4.0

- Disabled and Not Present are not available as link status indicators.
- Chassis is available as a link status indicator. If chassis is listed, the link is connected to a Summit Virtual Chassis.

Example

The following command displays port statistics for ports 1 through 3 on a stand-alone switch:

```
show ports 1-3 stats
```

The following command displays port statistics for slot 1, ports 1 through 3 on a modular switch:

```
show ports 1:1-1:3 stats
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 4.0 to support modular switches.

This command was modified in Extreme 4.1 to discontinue support for the chassis link status indicator.

Platform Availability

This command is available on all platforms.

show ports txerrors

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

Description

Displays real-time transmit error statistics.

For PoS modules, displays the `txerror` information for the PoS ports.

Syntax Description

portlist	Specifies a list of ports or slots and ports to which the parameters apply. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

N / A.

Usage Guidelines

If you do not specify a port number or range of ports, error statistics are displayed for all ports.

This status information may be useful for your technical support representative if you have a network problem.

For PoS modules, displays the `txerror` information for the PoS ports. Only a subset of the statistics displayed by this command are applicable to PoS ports. The fields that do not apply to PoS ports are displayed with values of all zeroes.

The following port transmit error information is collected by the switch:

- Port Number
- Link Status — The current status of the link. Options are:
  - Ready (R): The port is ready to accept a link.
  - Active (A): The link is present at this port.
  - Disabled (D): The link is disabled at this port.
  - Not Present (NP): The link is not 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 Lost Frames (TX Lost) — The total number of frames transmitted by the port that were lost.
- Transmit Parity Frames (TX Parity) — The bit summation has a parity mismatch.

For version 2.0 and 4.0

- Disabled and Not Present are not available as link status indicators.

Example

The following command displays transmit error statistics for ports 1 through 3 on a stand-alone switch:

```
show ports 1-3 txerrors
```

The following command displays transmit error statistics for slot 1, ports 1 through 3 on a modular switch:

```
show ports 1:1-1:3 txerrors
```

The output produced by the `show ports txerrors` command is similar to the following:

```
Port Tx Error Monitor                       Thu Dec 27 19:19:07 2001
Port         Link    Tx    Tx       Tx       Tx    Tx    Tx
            Status   Coll  Late Coll Deferred  Error  Lost  Parity
================================================================================
  1         A       0      0       0       0      0      0
  2         R       0      0       0       0      0      0
  3         R       0      0       0       0      0      0
  4         R       0      0       0       0      0      0
  5         R       0      0       0       0      0      0
  6         R       0      0       0       0      0      0
  7         R       0      0       0       0      0      0
  8         R       0      0       0       0      0      0



================================================================================
         Link Status: A-Active R-Ready D-Disabled NP-Not Present
           0->Clear Counters  U->page up  D->page down ESC->exit
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 6.1 to support PoS modules.

This command was modified in ExtremeWare 6.0 to support the Disabled and Not Present link status indicators.

This command was modified in ExtremeWare 4.0 to support modular switches.

Platform Availability

This command is available on all platforms.

show version

```
show version {detail}
```

Description

Displays the hardware serial numbers and versions, and software versions currently running on the switch.

Syntax Description

detail	Specifies display of slot board name and chassis or platform name.

Default

N/A.

Usage Guidelines

On chassis-based switches, displays the switch serial number and version numbers of MSM modules (BlackDiamond) and I/O modules (BlackDiamond and Alpine switches).

For PoS and MPLS modules, displays information that includes data about the PoS or MPLS module and the BootROM version of the PoS or MPLS module.

The following is an example of the type of information displayed when you execute the `show version` command:

- System Serial Number—A collection of numbers and letters that make up the serial number of the switch.
- CPU Serial Number—A collection of numbers and letters that make up the serial number of the CPU running in the switch. A rev number may also be listed.
- Image—The ExtremeWare software version currently running on the switch. If you have two software images downloaded on the switch, only the currently running ExtremeWare version information is displayed. The information displayed includes the version number, build number, and the software build date.
- BootROM—The BootROM version currently running on the switch.

If you use the `detail` option (supported in ExtremeWare version 6.2.1 or later) you may also see the following:

- Board/Chassis/Platform Name—The name of the system or module, inserted before the Serial Number in the display.

Depending on the model of your switch, the software running on your switch, and whether you have a stackable or modular switch, different version information may be displayed.

For PoS and MPLS modules:

The ExtremeWare technology release that supports the PoS or MPLS module includes multiple software packages. One software package runs on the MSM module while another package runs on each PoS or MPLS module. You must download the software packages independently using the ExtremeWare

download image command. Each software package has an associated version number that you can display using the show version command. It is recommended (not required), that the MSM software package and the MPLS module software package be the same version. To ensure compatibility, the MSM performs an automatic compatibility check before the PoS or MPLS module is activated. If the versions of the software packages are incompatible, the PoS or MPLS ports on the module will not come up and the show slot command will indicate that the software on the MPLS module is incompatible with the MSM software.

Example

The following command displays the hardware and software versions currently running on the switch:

```
show version
```

On a stackable switch, this command produces output similar to the following:

```
System Serial Number: 800078-11-0035M02442
CPU Serial Number: 700027-11  0034M-01445 CPLD Rev 04
Daughtercard Serial Number: 703015-02  0029M-02701 CPLD Rev ÿ
Image  :  Extremeware Version 6.2.0 (Build 60)  by Release_Master 09/21/0120:53:17
```

On a Black Diamond switch, this command produces output similar to the following:

```
Chassis:  801000-07-9946F00987
MSM A  :
MSM B  :  701021-08-0023F25758
SLOT 1 :  701026-03-0003Y00043
SLOT 2 :  701024-04-9949Y00055
SLOT 3 :  701005-09-9946F25172
SLOT 4 :
SLOT 5 :
SLOT 6 :  701028-01-0004Y00038
SLOT 7 :
SLOT 8 :

Image  :  Extremeware Version 6.2.0 (Build 60)  by Release_Master 09/21/0120:53:17

BootROM : 7.2
```

Using the detail option in the show version command produces output similar to the following on a Black Diamond switch:

```
Chassis  :      MSM64  801000-07-9946F00987
MSM    A :     MSM64i
MSM    B :     MSM64i  701021-08-0023F25758
SLOT   1 :      F48Ti  701026-03-0003Y00043
SLOT   2 :       G8Xi  701024-04-9949Y00055
SLOT   3 :       F32T  701005-09-9946F25172
SLOT   4 :      Empty
SLOT   5 :      Empty
SLOT   6 :       G8Ti  701028-01-0004Y00038
SLOT   7 :      Empty
SLOT   8 :      Empty


Image : Extremeware Version 6.2.1 (Build 18)  by Release_Master 02/14/02 15:04:26

BootROM : 7.2
```

History

This command was first available in ExtremeWare 2.0.

This command was available in an ExtremeWare IP Technology Services Release based on v6.1.8b12 to support MPLS modules.

This command was modified in ExtremeWare 6.1 to support PoS modules.

This command was modified in ExtremeWare 6.2.1 to support the `detail` option.

Platform Availability

This command is available on all platforms.

unconfig flowstats filter

```
unconfig flowstats filter <filter#>
```

Description

Removes the filter specification for the specified filter.

Syntax Description

filter#	Specifies the filter specification that should be removed.

Default

N/A.

Usage Guidelines

By unconfiguring the filter specification, this effectively disables this filter on all ports for which it was configured.

Example

The following command resets the values for filter 4:

```
unconfig flowstats filter 4
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

unconfig flowstats ports

```
unconfig flowstats ports <portlist>
```

Description

Resets the flow statistics configuration parameters for the specified ports to their default values.

Syntax Description

portlist	Specifies a set of ports or slots and ports that should be reset. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

N / A.

Usage Guidelines

This command does not affect the enabled or disabled status of flow collection on these ports, nor does it affect the configured export destinations.

Example

The following command resets the flow statistics configuration parameters for port 1 of slot 8 to their default values:

```
unconfig flowstats ports 8:1
```

History

This command was first available in ExtremeWare 6.1.5b20 for the PoS module

This command was made available in ExtremeWare 6.2 for all “i” series platforms.

Platform Availability

This command is available on the “*i*” series platforms.

13 STP Commands

This chapter describes:

- Commands related to creating, configuring, enabling, and disabling Spanning Tree Protocol (STP) on the switch
- Commands related to displaying and resetting STP settings on the switch

The Spanning Tree Protocol (STP) is a bridge-based mechanism for providing fault tolerance on networks. 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.

STP allows you to implement parallel paths for network traffic, and ensure that:

- Redundant paths are disabled when the main paths are operational.
- A redundant path is 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. After an STPD is created, one or more VLANs can be assigned to it.

A port can belong to multiple STPDs. In addition, a VLAN can span multiple STPDs.

The key points to remember when configuring VLANs and STP are:

- Each VLAN forms an independent broadcast domain.
- STP blocks paths to create a loop-free environment.
- When STP blocks a path, no data can be transmitted or received on the blocked port.
- Within any given STPD, all VLANs belonging to it use the same spanning tree.

If you delete a STPD, the VLANs that were members of that STPD are also deleted. You must remove all VLANs associated with the STP before deleting the STPD.

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.

Port Modes

An STP port has three modes of operation:

- 802.1D mode

  This mode is used for backward compatibility with previous STP versions and for compatibility with third-party switches using IEEE standard 802.1D. BPDUs are sent untagged in 1D mode. Because of this, on any given physical interface there can be only *one* STPD running in 1D mode.

- Extreme Multiple Instance Spanning Tree Protocol (EMISTP) mode

  EMISTP mode is an extension of STP that allows a physical port to belong to multiple STPDs by assigning the port to multiple VLANs. EMISTP adds significant flexibility to STP network design. BPDUs are sent with an 802.1Q tag having an STPD instance Identifier (StpdID) in the VLANid field.

- PVST+ mode

  This mode implements PVST+ in compatibility with third-party switches running this version of STP. The STPDs running in this mode have a one-to-one relationship with VLANs, and send and process packets in PVST+ format.

These port modes are for STP ports, not for physical ports. When a physical port belongs to multiple STPDs, it is associated with multiple STP ports It is possible for the physical port to run in different modes for different domains to which it belongs.

config stpd add vlan

```
config stpd <spanning tree name> add vlan <vlan name> {ports <portlist>
[dot1d | emistp | pvst-plus]
```

Description

Adds one or more VLANs, or a list of ports within a VLAN, to a specified STPD.

Syntax Description

spanning tree name	Specifies an STPD name on the switch.
vlan name	Specifies a VLAN name.
ports	Specifies the port or ports to be included in the STPD. (6.2)
dot1d	Specifies the STP port mode of operation to be 1D. (6.2)
emistp	Specifies the STP port mode of operation to be EMISTP. (6.2)
pvst-plus	Specifies the STP port mode of operation to be PVST+. (6.2)

Default

For versions up to 6.1, the default is N/A.

For version 6.2, all ports are in `emistp` mode, except those in STPD s0, whose default setting is `dot1d` mode.

Usage Guidelines

For version 6.2, this command adds a list of ports within a VLAN to a specified STPD. If no ports are specified, the entire VLAN is added.

For versions up to 6.1, this command adds one or more VLANs to the STPD. All VLANs participating in the STPD elect a Root Bridge and create a loop free least-cost path to the bridge.

Care must be taken to ensure that ports in overlapping domains do not interfere with the orderly working of each domain's protocol.

You must create a VLAN to add a VLAN to the STPD. To create a VLAN, use the `create vlan <name>` command.

You can create STP domains using the `create stpd <spanning tree name>` command.

For version 6.2:

Added keywords `dot1d`, `emistp`, and `pvst-plus` to specify STP port modes.

- `dot1d`—This mode is reserved for backward compatibility with previous STP versions. BPDUs are sent untagged in 1D mode. Because of this, on any given physical interface there can be only *one* STPD running in 1D mode.
- `emistp`—This mode sends BPDUs with an 802.1Q tag having an STPD instance Identifier (StpdID) in the VLANid field.

- pvst-plus—This mode implements PVST+ in compatibility with third-party switches running this version of STP. The STPDs running in this mode have a one-to-one relationship with VLANs, and send and process packets in PVST+ format.

An StpdID is used to identify each STP domain. You assign the StpdID when configuring the domain, and that VLAN cannot belong to another STPD. An StpdID must be identical to the VLANid of one of the member VLANs in that STP domain.

NOTE

These port modes are for STP ports, not for physical ports. When a physical port belongs to multiple STPDs, it is associated with multiple STP ports. It is possible for the physical port to run in different modes for different domains for which it belongs.

When the switch boots, it automatically creates a VLAN named *default* with a tag value of 1, and STPD s0 with an StpdID of 1. The switch associates VLAN *default* to STPD S0. By default, all ports that belong to this VLAN and STPD in dot1D mode.

Example

Create a VLAN named *marketing* and an STPD named *STPD1* as follows:

```
create vlan marketing
create stpd stpd1
```

The following command adds the VLAN named *marketing* to the STPD *STPD1:*

```
config stpd stpd1 add vlan marketing
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 6.2 to support STP port mode configurations.

Platform Availability

This command is available on all platforms.

config stpd delete vlan

```
config stpd <spanning tree name> delete vlan <vlan name> {ports <portlist>
```

Description

Deletes one or more ports in the specified VLAN from an STPD.

Syntax Description

spanning tree name	Specifies an STPD name on the switch.
vlan name	Specifies a VLAN name.
ports	Specifies the port or ports to be included in the STPD.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes a VLAN named *Marketing* from the STPD *STPD1:*

```
config stpd stpd1 delete vlan marketing
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all platforms.

config stpd forwarddelay

```
config stpd <spanning tree name> forwarddelay <seconds>
```

Description

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.

Syntax Description

spanning tree name	Specifies an STPD name on the switch.
seconds	Specifies the forward delay time in seconds.

Default

15 seconds.

Usage Guidelines

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 range for the <seconds> parameter is 4 through 30 seconds.

Example

The following command sets the forward delay from *STPD1* to 20 seconds:

```
config stpd stpd1 forwarddelay 20
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config stpd hellotime

```
config stpd <spanning tree name> hellotime <seconds>
```

Description

Specifies the time delay (in seconds) between the transmission of Bridge Protocol Data Units (BPDUs) from this STPD when it is the Root Bridge.

Syntax Description

spanning tree name	Specifies an STPD name on the switch.
seconds	Specifies the hello time in seconds.

Default

2 seconds.

Usage Guidelines

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 range for the <seconds> parameter is 1 through 10 seconds.

Example

The following command sets the time delay from *STPD1* to 10 seconds:

```
config stpd stpd1 hellotime 10
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config stpd maxage

```
config stpd <spanning tree name> maxage <seconds>
```

Description

Specifies the maximum age of a BPDU in the specified STPD.

Syntax Description

spanning tree name	Specifies an STPD name on the switch.
seconds	Specifies the maxage time in seconds.

Default

20 seconds.

Usage Guidelines

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 range for the <seconds> parameter is 6 through 40 seconds.

Note that the time must be greater than, or equal to 2 * (Hello Time + 1) and less than, or equal to 2 * (Forward Delay –1).

Example

The following command sets the maximum age of *STPD1* to 30 seconds:

```
config stpd stpd1 maxage 30
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config stpd ports cost

```
config stpd <spanning tree name> ports cost <cost> [<portlist>]
```

Description

Specifies the path cost of the port in the specified STPD.

Syntax Description

spanning tree name	Specifies an STPD name on the switch.
cost	Specifies a numerical port cost value.
portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

The switch automatically assigns a default path cost based on the speed of the port, as follows:

- For a 10Mbps port, the default cost is 100.
- For a 100Mbps port, the default cost is 19.
- For a 1000Mbps port, the default cost is 4.

Usage Guidelines

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.

On a modular switch, `<portlist>` can be a list of slots and ports. On a stand-alone switch, `<portlist>` can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

The range for the `cost` parameter is 1 through 65,535. The switch automatically assigns a default path cost based on the speed of the port.

Example

The following command configures a cost of 100 to ports 1 through 5 in STPD *s0* on a stand-alone switch:

```
config stpd s0 ports cost 100 1-5
```

The following command configures a cost of 100 to slot 2, ports 1 through 5 in STPD *s0* on a modular switch:

```
config stpd s0 ports cost 100 2:1-2:5
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 4.0 to support modular switches.

Platform Availability

This command is available on all platforms.

config stpd ports mode

```
config stpd <spanning tree name> ports mode {dot1d | emistp | pvst-plus}
[<portlist>]
```

Description

Configures the STP mode of operation for the specified port list.

Syntax Description

spanning tree name	Specifies an STPD name on the switch.
dot1d	Specifies IEEE 802.1d-compliant packet formatting. A physical port can only be a member of one STPD running it `dot1d` mode.
emistp	Specifies 802.1d formatting and 802.1q tagging.
pvst-plus	Specifies PVST+ packet formatting.
portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

Ports in the default STPD (s0) are `dot1d` mode. Ports in user-created STPDs are in `emistp` mode.

Usage Guidelines

None.

Example

The following command configures STPD *s1* with PVST+ packet formatting for port 1 on a stand-alone switch:

```
config stpd s1 ports mode pvst-plus 1
```

The following command configures STPD *s1* with PVST+ packet formatting for slot 2, port 1 on a modular switch:

```
config stpd s1 ports mode pvst-plus 2:1
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all platforms.

config stpd ports priority

```
config stpd <spanning tree name> ports priority <priority> [<portlist>]
```

Description

Specifies the port priority of the port in the specified STPD.

Syntax Description

spanning tree name	Specifies an STPD name on the switch.
priority	Specifies a numerical port priority value.
portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

For version 6.0 and higher, the default setting is 16.

For version 2.0 and 4.0, the default setting is 128.

Usage Guidelines

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.

By changing the priority of the port, you can make it more or less likely to become the root port or a designated port.

A setting of 0 indicates the lowest priority.

On a modular switch, `<portlist>` can be a list of slots and ports. On a stand-alone switch, `<portlist>` can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

For version 6.0 and higher:

- The range for the `priority` parameter is 0 through 31.

For version 2.0 and 4.0:

- The range for the `priority` parameter is 0 through 255.

Example

The following command assigns a priority of 1 to ports 1 through 5 in STPD *s0* on a stand-alone switch:

```
config stpd s0 ports priority 1 1-5
```

The following command assigns a priority of 1 to slot 2, ports 1 through 5 in STPD *s0* on a modular switch:

```
config stpd s0 ports priority 1 2:1-2:5
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 6.1 to update the `priority` parameter.

This command was modified in ExtremeWare 4.0 to support modular switches.

Platform Availability

This command is available on all platforms.

config stpd priority

```
config stpd <spanning tree name> priority <priority>
```

Description

Specifies the bridge priority of the STPD.

Syntax Description

spanning tree name	Specifies an STPD name on the switch.
priority	Specifies the bridge priority of the STPD.

Default

32,768.

Usage Guidelines

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.

By changing the priority of the STPD, you can make it more or less likely to become the root bridge.

The range for the `priority` parameter is 0 through 65,535. A setting of 0 indicates the highest priority.

Example

The following command sets the bridge priority of *STPD1* to 16,384:

```
config stpd stpd1 priority 16384
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config stpd tag

```
config stpd <spanning tree name> tag <vlan tag>
```

Description

Assigns an StpdID to an STPD.

Syntax Description

spanning tree name	Specifies an STPD name on the switch.
vlan tag	Specifies the VLANid of a VLAN that is owned by the STPD.

Default

N/A.

Usage Guidelines

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.

An StpdID is used to identify each STP domain. You assign the StpdID when configuring the domain, and that VLAN cannot belong to another STPD. An StpdID must be identical to the VLANid of one of the member VLANs in that STP domain. Unless all ports are running in 802.1d mode, an STPD must be configured with an StpdID.

You must create and configure the VLAN, along with the tag, before you can configure the STPD tag. To create a VLAN, use the `create vlan` command. To configure the VLAN, use the `config vlan` command.

In addition to the VLAN attributes that you will use in the STPD, you must first create an STPD. To create an STPD, use the `create stpd` command.

Example

The following command assigns an StpdID to the `purple_st` STPD:

```
config stpd purple_st tag 200
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all platforms.

config vlan add ports stpd

```
config vlan <vlan name> add ports <portlist> stpd <spanning tree name>
[dot1d | emistp | pvst-plus]
```

Description

Adds a list of ports within a VLAN to a specified STPD.

Syntax Description

vlan name	Specifies a VLAN name.
portlist	Specifies the port or ports to be included in the STPD.
spanning tree name	Specifies an STPD name on the switch.
dot1d	Specifies the STP port mode of operation to be 1D.
emistp	Specifies the STP port mode of operation to be EMISTP.
pvst-plus	Specifies the STP port mode of operation to be PVST+.

Default

All ports are in `emistp` mode, except those in STPD s0, whose default setting is `dot1d` mode.

Usage Guidelines

This command performs the same function as the `config stpd add vlan` command with the ports option included.

This command adds a list of ports within a VLAN to a specified STPD, and specifies the mode for those ports.

- `dot1d`—In this mode, BPDUs are sent untagged in 1D mode. Because of this, on any given physical interface there can be only *one* STPD running in 1D mode. This mode supports the industry standard implementation, and can be used with non-Extreme devices. It can also be used for backward compatibility with previous STP versions.
- `emistp`—This mode sends BPDUs with an 802.1Q tag having an STPD instance Identifier (StpdID) in the VLANid field. This is an Extreme proprietary mode, and cannot be used with non-Extreme devices.
- `pvst-plus`—This mode implements PVST+ in compatibility with third-party switches running this version of STP. The STPDs running in this mode have a one-to-one relationship with VLANs, and send and process packets in PVST+ format.

These port modes are for STP ports, not for physical ports. When a physical ports belongs to multiple STPDs, it is associated with multiple STP ports. It is possible for the physical port to run in different modes for different domains for which it belongs.

Example

The following command adds ports 2 and 3, members of a VLAN named *Marketing*, to the STPD named *STPD1, and specifies that they be in EMISTP mode:*

```
config vlan marketing add ports 2-3 stpd stpd1 emistp
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all "*i*"-series platforms.

create stpd

```
create stpd <name>
```

Description

Creates a user-defined STPD.

Syntax Description

name	Specifies a user-defined STPD name.

Default

The default device configuration contains a single STPD called *s0*.

When an STPD is created, the STPD has the following default parameters:

- State — disabled
- StpdID — none
- Assigned VLANs — none
- Bridge priority — 32,768
- Hello time — 2 seconds
- Forward delay — 15 seconds
- Port mode — Ports in the default STPD (s0) are `dot1d` mode. Ports in user-created STPDs are in `emistp` mode.

Usage Guidelines

Each STPD name must be unique, and cannot duplicate any other named elements on the switch (such as VLANs, QoS profiles, Access profiles, or route maps). If you are uncertain about the VLAN profile names on the switch, use the `show vlan` command to view the VLAN profiles. If you are uncertain about QoS profile names on the switch, use the `show qos <qos profile>` command to view the QoS profiles.

Each STPD has its own Root Bridge and active path. After the STPD is created, one or more VLANs can be assigned to it.

Example

The following example creates an STPD named *purple_st*:

```
create stpd purple_st
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

delete stpd

```
delete stpd <spanning tree name>
```

Description

Removes a user-defined STPD from the switch.

Syntax Description

spanning tree name	Specifies a user-defined STPD name on the switch.

Default

N / A.

Usage Guidelines

If you remove an STPD, the VLANs that were members of that STPD are also deleted. An STPD can only be removed if all VLANs have been deleted from it.

The default STPD, *s0*, cannot be deleted.

Example

The following command deletes an STPD named *purple_st*:

```
delete stpd purple_st
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

disable ignore-bpdu vlan

```
disable ignore-bpdu vlan <name>
```

Description

Allows the switch to recognize STP BDUs.

Syntax Description

name	Specifies a VLAN name.

Default

Disabled.

Usage Guidelines

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. 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.

Example

The following command disables the ignore-stp option on the VLAN *accounting*:

```
disable ignore-stp accounting
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

disable ignore-stp vlan

```
disable ignore-stp vlan <name>
```

Description

Allows a VLAN to use STP port information.

Syntax Description

name	Specifies a VLAN name.

Default

Disabled.

Usage Guidelines

The `vlan` keyword is optional.

Example

The following command disables the ignore-stp option on the VLAN *accounting*:

```
disable ignore-stp accounting
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all platforms.

disable stpd

```
disable stpd {<spanning tree name>}
```

Description

Disables the STP mechanism on a particular STPD or for all STPDs.

Syntax Description

spanning tree name	Specifies an STPD name on the switch.

Default

Disabled.

Usage Guidelines

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

The `spanning tree name` keyword is optional. You do not need to indicate an STPD name if you disable the STP mechanism for all STPDs.

Example

The following command disables an STPD named *purple_st*:

```
disable stpd purple_st
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

disable stpd ports

```
disable stpd <spanning tree name> ports [<portlist>]
```

Description

Disables STP on one or more ports.

Syntax Description

spanning tree name	Specifies an STPD name on the switch.
portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

N / A.

Usage Guidelines

Disabling STP on one or more ports puts those ports in *forwarding* state; all Bridge Protocol Data Units (BPDUs) received on those ports will be disregarded and dropped.

On a modular switch, `<portlist>` can be a list of slots and ports. On a stand-alone switch, `<portlist>` can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

You must create one or more STP domains, configure, and enable an STPD before you can use the `disable stpd port` command.

Example

The following command disables port 4 on an STPD named *Backbone_st* on a stand-alone switch:

```
disable stpd backbone_st ports 4
```

The following command disables slot 2, port 4 on an STPD named *Backbone_st* on a modular switch:

```
disable stpd backbone_st ports 2:4
```

History

This command was first available in ExtremeWare 2.0.

This command was modified in ExtremeWare 4.0 to support modular devices.

Platform Availability

This command is available on all platforms.

disable stpd rapid-root-failover

```
disable stpd <spanning tree name> rapid-root-failover
```

Description

Disables rapid root failover for STP recovery times.

Syntax Description

spanning tree name	Specifies an STPD name on the switch.

Default

Disabled.

Usage Guidelines

To view the status of rapid root failover on the switch, use the `show stpd` command. The `show stpd` command displays information about the STPD configuration on the switch including the enable / disable state for rapid root failover.

Example

The following command disables rapid root fail over on STPD *Backbone_st*:

```
disable stpd backbone_st rapid-root-failover
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all platforms.

enable ignore-bpdu vlan

```
enable ignore-bpdu vlan <vlan name>
```

Description

Configures the switch to ignore the STP BPDUs, which prevents ports in the VLAN from becoming part of an STPD.

Syntax Description

vlan name	Specifies a VLAN name.

Default

Disabled.

Usage Guidelines

This command is useful when you have a known topology with switches outside your network, and you wish to keep the root bridge within your network.

Example

The following command configures the switch to ignore STP BPDUs t on the VLAN *accounting*:

```
enable ignore-bpdu vlan accounting
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all platforms.

enable ignore-stp vlan

```
enable ignore-stp vlan <vlan name>
```

Description

Configures the switch to ignore the STP protocol and not block traffic for the VLAN(s).

Syntax Description

vlan name	Specifies a VLAN name.

Default

Disabled.

Usage Guidelines

This command is useful when multiple VLANs share the same physical ports, but only some of the VLANs require STP protection.

Example

The following command enables the ignore-stp option on the VLAN *accounting*:

```
enable ignore-stp accounting
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

enable stpd

```
enable stpd {<spanning tree name>}
```

Description

Enables the STP protocol for one or all STPDs.

Syntax Description

spanning tree name	Specifies an STPD name on the switch.

Default

Disabled.

Usage Guidelines

You must create one or more STP domains and configure an STPD before you can use the `enable stpd` command. Use the `create stpd <spanning tree name>` command to create an STPD.

Example

The following command enables an STPD named *Backbone_st*:

```
enable stpd backbone_st
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

enable stpd rapid-root-failover

```
enable stpd <spanning tree name> rapid-root-failover
```

Description

Enables rapid root failover for faster STP recovery times.

Syntax Description

spanning tree name	Specifies an STPD name on the switch.

Default

Disabled.

Usage Guidelines

To view the status of rapid root failover on the switch, use the `show stpd` command. The `show stpd` command displays information about the STPD configuration on the switch including the enable/disable state for rapid root failover.

Example

The following command enables rapid root fail over on STPD *Backbone_st*:

```
enable stpd backbone_st rapid-root-failover
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all platforms.

enable stpd ports

```
enable stpd ports <portlist>
```

Description

Enables the STP protocol on one or more ports.

Syntax Description

portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.

Default

Enabled.

Usage Guidelines

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.

You must create and configure one or more STP domains before you can use the `enable stpd ports` command. Use the `create stpd <name>` command to create an STP domain. If you have considerable knowledge and experience with STP, you can configure the STPD using the `config stpd` commands. However, the default STP parameters are adequate for most networks.

On a modular switch, `<portlist>` can be a list of slots and ports. On a stand-alone switch, `<portlist>` can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

Example

The following command enables port 4 on an STPD named *Backbone_st* on a stand-alone switch:

```
enable stpd backbone_st ports 4
```

The following command enables slot 2, port 4 on an STPD named *Backbone_st* on a modular switch:

```
enable stpd backbone_st ports 2:4
```

History

This command was first available in ExtremeWare 2.0.

Support for modular switches was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

show stpd

```
show stpd {detail | <spanning tree name>}
```

Description

Displays STPD settings on the switch.

Syntax Description

detail	Specifies that STPD settings should be shown for each STPD. (Version 6.2)
spanning tree name	Specifies an STPD on the switch.

Default

N / A.

Usage Guidelines

The command displays the following STPD information:

- STPD name
- Bridge ID
- STPD configuration information

You must create, configure, and enable one or more STP domains before you can use the `show stpd` command. Use the `create stpd <name>` command to create an STP domain. If you have considerable knowledge and experience with STP, you can configure the STPD using the `config stpd` commands. However, the default STP parameters are adequate for most networks. Use the `create stpd <name>` command to create an STPD. Use the `enable stpd {<spanning tree name>}` command to enable an STPD.

Example

The following command displays STPD settings on an STPD named *Backbone_st:*

```
show stpd backbone_st
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

show stpd ports

```
show stpd <spanning tree name> ports [<portlist> | all] {detail}
```

Description

Displays the STP state of a port.

Syntax Description

spanning tree name	Specifies an STPD name.
portlist	Specifies one or more ports or slots and ports. On a modular switch, can be a list of slots and ports. On a stand-alone switch, can be one or more port numbers. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
all	Specifies that STPD state information should be displayed for all ports. (versions 4.0, 6.0)
detail	Specifies that STPD state information should be displayed for all ports, or for the ports in the port list. (version 6.2 and higher)

Default

N / A.

Usage Guidelines

This command displays the following:

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

On a modular switch, `<portlist>` can be a list of slots and ports. On a stand-alone switch, `<portlist>` can be one or more port numbers. For a detailed explanation of port specification, see "Modular Switch Numerical Ranges" or "Stand-alone Switch Numerical Ranges" in Chapter 1.

For version 6.2 and higher:

- When used without a port list, the `detail` option replaces the `all` option to specify that information should be displayed for all ports.

Example

The following command displays the state of port 4 on an STPD named *Backbone_st* on a stand-alone switch:

```
show stpd ports Backbone_st 4
```

The following command displays the state of slot 2, port 4 on an STPD named *Backbone_st* on a modular switch:

```
show stpd ports Backbone_st 2:4
```

History

This command was first available in ExtremeWare 2.0.

Support for modular switches was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

unconfig stpd

```
unconfig stpd {<spanning tree name> | all}
```

Description

Restores default STP values to a particular STPD or all STPDs.

Syntax Description

spanning tree name	Specifies an STPD name on the switch.
all	Specifies all STPDs on the switch. (2.0)

Default

N/A.

Usage Guidelines

Use this command to restore default STP values to a particular STPD.

For version 2.0:

- You can use the `all` parameter to specify all STPDs.

Example

The following command restores default values to an STPD named *Backbone_st*:

```
unconfig stpd backbone_st
```

History

This command was first available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

14 ESRP Commands

This chapter describes the following commands:

- Commands for enabling and disabling ESRP
- Commands for performing basic ESRP configuration
- Commands for enabling and disabling port restart and failure tracking for ESRP and VRRP

ESRP is a feature of ExtremeWare that allows multiple switches to provide redundant layer 3 routing services to users. In addition to providing layer 3 routing redundancy, ESRP also provides for layer 2 redundancy. These "layered" redundancy features can be used in combination or independently. The layer 2 redundancy features of ESRP offer fast failure recovery and provide for dual-homed system design. In some instances, depending on network system design, ESRP can provide better resiliency than using the Spanning Tree Protocol (STP).

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

ESRP is configured on a per-VLAN basis on each switch. A maximum of four switches can participate in providing redundant layer 3 or layer 2 services to a single VLAN. A maximum of 3000 VLANs can run ESRP simultaneously on a single switch. The switches exchange keep-alive packets for each VLAN independently. Only one switch can actively provide layer 3 routing and / or layer 2 switching for each VLAN. The switch performing the forwarding for a particular VLAN is considered the "master" for that VLAN. Other participating switches for the VLAN are in standby mode.

To have two or more switches participate in ESRP, the following must be true:

- For each VLAN to be made redundant, the switches must have the ability to exchange packets on the same layer 2 broadcast domain for that VLAN. Multiple paths of exchange can be used.
- For a VLAN to be recognized as participating in ESRP, the assigned IP address or the IPX NETid for the separate switches must be *identical*. Other aspects of the VLAN, including its name, are ignored.
- ESRP must be enabled on the desired VLANs for each switch. ESRP cannot be enabled on the VLAN "default."
- Extreme Discovery Protocol (EDP) must be enabled on the ports that are members of the ESRP VLANs. (The default setting is enabled.)

config esrp port-mode ports

```
config esrp port-mode [host | normal] ports <portlist> {don't-count}
```

Description

Configures the ESRP port mode for ESRP host attach.

Syntax Description

host	Specifies that the ports should be configured as host ports.
normal	Specifies that the ports should be configured as normal ports.
portlist	Specifies the list of ports that should be configured. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
don't-count	Specifies that host ports should not be counted as active ports.

Default

Normal.

Usage Guidelines

This feature is useful in dual-homed server environments in conjunction with high availability server load-balancing (SLB) configurations.

Ports configured as normal ports do not accept or transmit Layer 2 or Layer 3 traffic when the local ESRP device is a slave.

Ports configured as host ports allow configured ports that do not represent loops to the network to continue operation independent of ESRP status.

`don't-count` has the effect of not counting the host ports as active ports. This has the convenience of minimal ESRP state changes due to frequent client activities like reboots and unplugging laptops.

An L2 connection for VLANs between ESRP switches is required.

Example

The following command configures ports 1 through 5 as host ports, and prevents them from being counted as active ports:

```
config esrp port-mode host ports 1-5 don't-count
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan add domain-member vlan

```
config vlan <super_esrp_vlan> add domain-member vlan <sub_esrp_vlan>
```

Description

Adds a VLAN to an ESRP domain.

Syntax Description

super_esrp_vlan	Specifies the name of an ESRP-enabled super-VLAN.
sub_esrp_vlan	Specifies the name of a sub-VLAN.

Default

N/A.

Usage Guidelines

One port of each member VLAN must belong to the domain master VLAN.

ESRP is performed in the domain master VLAN only, and not the other domain members.

Example

The following command adds the sub-VLAN *sub_esrp1* to ESRP-enabled super VLAN *esrp-super*:

```
config vlan esrp-super add domain-member vlan sub_esrp1
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan add ports no-restart

```
config vlan <name> add ports [<portlist> | all] no-restart
```

Description

Disables port restart for a port.

Syntax Description

name	Specifies an ESRP-enabled or VRRP-enabled VLAN name.
portlist	Specifies list of ports or slots and ports. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8.
	`all` indicates that all ports should have restart disabled.

Default

N/A.

Usage Guidelines

To disable port restart, you must delete the ports, and then add them again with the no-restart option.

This command can be used with both ESRP-enabled VLANs and VRRP VLANs.

Example

The following two commands disable port restart for ports 7-9 on VLAN esrp1:

```
config vlan esrp1 delete ports 7-9
config vlan esrp1 add ports 7-9 no-restart
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan add ports restart

```
config vlan <name> add ports [<portlist> | all] restart
```

Description

Configures ESRP or VRRP to restart ports if those ports are members of a VLAN that becomes a slave.

Syntax Description

name	Specifies an ESRP-enabled or VRRP VLAN name.
portlist	Specifies list of ports or slots and ports. May be in the form 1, 2, 3-5, 2:5, 2:6-2:8. `all` indicates that all ports should have restart enabled.

Default

N/A.

Usage Guidelines

If a VLAN becomes a slave, ESRP or VRRP disconnects member ports that have port restart enabled. The disconnection of these ports causes downstream devices to remove the ports from their FDB tables. After 3 seconds the ports re-establish connection with the ESRP- or VRRP-enabled device. This feature allows you to use ESRP or VRRP in networks that include equipment from other vendors.

This command can be used with both ESRP-enabled VLANs and VRRP VLANs.

Example

The following command enables port restart for ports 7-9 on VLAN *esrp1*:

```
config vlan esrp1 add ports 7-9 restart
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan add track-bgp

```
config vlan <name> add track-bgp failover <priority>
```

Description

Configures an ESRP-enabled VLAN to track any available BGP route.

Syntax Description

name	Specifies an ESRP-enabled VLAN name.
priority	Specifies a number between 0 and 254.

Default

No BGP route tracking.

Usage Guidelines

The switch cannot be the ESRP master if none of the specified routes are reachable.

If a failure is detected, the ESRP VLAN priority steps to the failover-priority value specified. By setting the failover priority to be lower than the normal priority of the VLAN, it will cause the affected VLAN to go into standby.

The range of the priority value is 0 to 254. Setting the priority to 255 configures the switch to standby mode, and makes it ineligible to become the master. The switch will remain in standby mode even when the VLAN fails over from the current master.

To make effective use of this feature, the following should be true:

- The priority field should be given precedence over the other election factors by assigning the `priority-ports-track-mac` election algorithm to the VLAN.
- The normal priority of the ESRP-enabled VLANs must be higher than the failover priority of this command.

Example

The following command enables BGP failure tracking, and specifies that the ESRP priority for VLAN esrp-1 be set to 10 when no BGP routes are reachable.

```
config vlan esrp-1 add track-bgp failover 10
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan add track-diagnostic

```
config vlan <name> add track-diagnostic failover <priority>
```

Description

Configures backplane diagnostics failure tracking for an ESRP-enabled VLAN.

Syntax Description

name	Specifies a VLAN name.
priority	Specifies a number between 0 and 255.

Default

No diagnostic tracking.

Usage Guidelines

If a diagnostic failure is detected, the ESRP VLAN priority steps to the failover-priority value specified. By setting the failover priority to be lower than the normal priority of the VLAN, it will cause the affected VLAN to go into standby.

The range of the priority value is 0 to 254. Setting the priority to 255 configures the switch to standby mode, and makes it ineligible to become the master. The switch will remain in standby mode even when the VLAN fails over from the current master.

To make effective use of this feature, the following should be true:

- The priority field should be given precedence over the other election factors by assigning the `priority-ports-track-mac` election algorithm to the VLAN.
- The normal priority of the ESRP-enabled VLANs must be higher than the failover priority of this command.

Example

The following command enables diagnostic failure tracking, and specifies that the ESRP priority for VLAN esrp-1 be set to 10 upon a diagnostic failure.

```
config vlan esrp-1 add track-diagnostic failover 10
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan add track-environment

```
config vlan <name> add track-environment failover <priority>
```

Description

Configures an ESRP-enabled VLAN to track environmental failures.

Syntax Description

name	Specifies a VLAN name.
priority	Specifies a number between 0 and 255.

Default

No environmental tracking.

Usage Guidelines

Environmental tracking tracks fan, power supply, and chassis temperature status.

If a failure is detected, the ESRP VLAN priority steps to the failover-priority value specified. By setting the failover priority to be lower than the normal priority of the VLAN, it will cause the affected VLAN to go into standby.

The range of the priority value is 0 to 254. Setting the priority to 255 configures the switch to standby mode, and to be ineligible to become the master. The switch will remain in standby mode even when the VLAN fails over from the current master.

To make effective use of this feature, the following should be true:

- The priority field should be given precedence over the other election factors by assigning the `priority-ports-track-mac` election algorithm to the VLAN.
- The normal priority of the ESRP-enabled VLANs must be higher than the failover priority of this command.

Example

The following command enables diagnostic failure tracking, and specifies that the ESRP priority for VLAN esrp-1 be set to 10 upon a diagnostic failure.

```
config vlan esrp-1 add track-environment failover 10
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan add track-iproute

```
config vlan <name> add track-iproute <ipaddress>/<masklength>
```

Description

Configures an ESRP-enabled VLAN or a VRRP VLAN to track a route entry in the kernel route table.

Syntax Description

name	Specifies an ESRP-enabled or VRRP VLAN name.
ipaddress	Specifies the IP address of the route entry to be tracked.

Default

No route tracking.

Usage Guidelines

If the specified routes are not reachable, the device automatically relinquishes master status and remains in standby mode (for ESRP) or backup mode (for VRRP).

This command can be used with both ESRP-enabled VLANs and VRRP VLANs.

Example

The following command enables IP route failure tracking for routes to the specified subnet:

```
config vlan esrp-1 add track-iproute 192.168.46.0/24
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan add track-ospf

```
config vlan <name> add track-ospf failover <priority>
```

Description

Configures an ESRP-enabled VLAN to track any available OSPF route.

Syntax Description

name	Specifies an ESRP-enabled VLAN name.
priority	Specifies a number between 0 and 255.

Default

No OSPF route tracking.

Usage Guidelines

The switch cannot be the ESRP master if none of the specified routes are reachable.

If a failure is detected, the ESRP VLAN priority steps to the failover-priority value specified. By setting the failover priority to be lower than the normal priority of the VLAN, it will cause the affected VLAN to go into standby.

The range of the priority value is 0 to 254. Setting the priority to 255 configures the switch to standby mode, and to be ineligible to become the master. The switch will remain in standby mode even when the VLAN fails over from the current master.

To make effective use of this feature, the following should be true:

- The priority field should be given precedence over the other election factors by assigning the `priority-ports-track-mac` election algorithm to the VLAN.
- The normal priority of the ESRP-enabled VLANs must be higher than the failover priority of this command.

Example

The following command enables OSPF route failure tracking, and specifies that the ESRP priority for VLAN esrp-1 be set to 10 when all OSPF routes become unreachable:

```
config vlan esrp-1 add track-ospf failover 10
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan add track-ping

```
config vlan <name> add track-ping <ipaddress> frequency <seconds> miss
<number>
```

Description

Configures an ESRP-enabled VLAN or VRRP VLAN to track an external gateway using ping.

Syntax Description

name	Specifies an ESRP-enabled or VRRP VLAN name.
ipaddress	Specifies the IP address of the external gateway.
seconds	Specifies the interval in seconds between ping requests.
number	Specifies the number of consecutive ping failures that will initiate failover to an ESRP slave or VRRP backup router.

Default

No ping tracking. Default miss number for VRRP is 3 consecutive missed ping responses.

Usage Guidelines

If the external gateway is not reachable as indicated by consecutive ping failures, the device automatically relinquishes master status and remains in standby mode (for ESRP) or backup mode (for VRRP).

This command can be used with both ESRP-enabled VLANs and VRRP VLANs.

Example

The following command enables ping tracking for the external gateway at 10.207.29.17, pinging every 10 seconds, and considering the gateway to be unreachable if no response is received to 5 consecutive pings:

```
config vlan esrp-1 add track-ping 10.207.29.17 frequency 10 miss 5
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan add track-rip

```
config vlan <name> add track-rip failover <priority>
```

Description

Configures an ESRP-enabled VLAN to track any available RIP route.

Syntax Description

name	Specifies an ESRP-enabled VLAN name.
priority	Specifies a number between 0 and 254.

Default

No RIP route tracking.

Usage Guidelines

The switch cannot be the ESRP master if none of the specified routes are reachable.

If a failure is detected, the ESRP VLAN priority steps to the failover-priority value specified. By setting the failover priority to be lower than the normal priority of the VLAN, it will cause the affected VLAN to go into standby.

The range of the priority value is 0 to 254. Setting the priority to 255 configures the switch to standby mode, and to be ineligible to become the master. The switch will remain in standby mode even when the VLAN fails over from the current master.

To make effective use of this feature, the following should be true:

- The priority field should be given precedence over the other election factors by assigning the `priority-ports-track-mac` election algorithm to the VLAN.
- The normal priority of the ESRP-enabled VLANs must be higher than the failover priority of this command.

Example

The following command enables RIP route tracking, and specifies that the ESRP priority for VLAN *esrp-1* be set to 10 upon a diagnostic failure:

```
config vlan esrp-1 add track-rip failover 10
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan add track-vlan

```
config vlan <name> add track-vlan <vlan_tracked>
```

Description

Configures an ESRP-enabled VLAN or a VRRP VLAN to track port connectivity to a specified VLAN.

Syntax Description

name	Specifies an ESRP-enabled or VRRP VLAN name.
vlan_tracked	Specifies the VLAN to be tracked.

Default

Disabled.

Usage Guidelines

If no active ports remain on the specified VLANs, the device automatically relinquishes master status and remains in standby mode (for ESRP) or backup mode (for VRRP).

An ESRP or VRRP VLAN can track one VLAN.

This command can be used with both ESRP-enabled VLANs and VRRP VLANs.

Example

The following command enables ESRP-enabled VLAN *esrp-1* to track port connectivity to VLAN *engineering*:

```
config vlan esrp-1 add track-vlan engineering
```

History

This command was first available in ExtremeWare 4.1.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan delete domain-member vlan

```
config vlan <super_esrp_vlan> delete domain-member vlan <sub_esrp_vlan>
```

Description

Deletes a VLAN from an ESRP domain.

Syntax Description

super_esrp_vlan	Specifies a VLAN name.
sub_esrp_vlan	Specifies a number in seconds.

Default

N/A.

Usage Guidelines

None.

Example

The following command adds the sub-VLAN *sub_esrp1* to ESRP-enabled super VLAN *esrp-super*:

```
config vlan esrp-super delete domain-member vlan sub_esrp1
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan delete track-bgp

```
config vlan <name> delete track-bgp
```

Description

Disables BGP route tracking for an ESRP-enabled VLAN.

Syntax Description

name	Specifies an ESRP-enabled VLAN name.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables diagnostic failure tracking for VLAN esrp-1:

```
config vlan esrp-1 delete track-bgp
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan delete track-diagnostic

```
config vlan <name> delete track-diagnostic
```

Description

Disables backplane diagnostics failure tracking for an ESRP-enabled VLAN.

Syntax Description

name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables diagnostic failure tracking for VLAN *esrp-1*:

```
config vlan esrp-1 delete track-diagnostic
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan delete track-environment

```
config vlan <name> delete track-environment
```

Description

Disables backplane environmental failure tracking.

Syntax Description

name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables environmental failure tracking for VLAN *esrp-1*:

```
config vlan esrp-1 delete track-environment
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan delete track-iproute

```
config vlan <name> delete track-iproute <ipaddress>/<masklength>
```

Description

Disables route table entry tracking for an ESRP-enabled VLAN or a VRRP VLAN.

Syntax Description

name	Specifies an ESRP-enabled or VRRP VLAN name.
ipaddress	Specifies the IP address of the route entry to be tracked.

Default

N / A.

Usage Guidelines

This command can be used with both ESRP-enabled VLANs and VRRP VLANs.

Example

The following command disables tacking of routes to the specified subnet for VLAN *esrp-1*:

```
config vlan esrp-1 add track-iproute 192.168.46.0/24
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan delete track-ospf

```
config vlan <name> delete track-ospf
```

Description

Disables OSPF route tracking for an ESRP-enabled VLAN.

Syntax Description

name	Specifies an ESRP-enabled VLAN name.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables BGP route tracking for VLAN *esrp-1*:

```
config vlan esrp-1 delete track-ospf
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan delete track-ping

```
config vlan <name> delete track-ping <ipaddress>
```

Description

Disables the tracking of an external gateway using ping.

Syntax Description

name	Specifies an ESRP-enabled or VRRP VLAN name.
ipaddress	Specifies the IP address of the external gateway.

Default

N / A.

Usage Guidelines

This command can be used with both ESRP-enabled VLANs and VRRP VLANs.

Example

The following command disables ping tracking for the external gateway at 10.207.29.17:

```
config vlan esrp-1 delete track-ping 10.207.29.17
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan delete track-rip

```
config vlan <name> delete track-rip
```

Description

Disables RIP route tracking for an ESRP-enabled VLAN.

Syntax Description

name	Specifies an ESRP-enabled VLAN name.

Default

No RIP route tracking.

Usage Guidelines

None.

Example

The following command disables RIP route failure tracking for VLAN *esrp-1*:

```
config vlan esrp-1 delete track-rip
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan delete track-vlan

```
config vlan <name> delete track-vlan <vlan_tracked>
```

Description

Disables the tracking of port connectivity to a specified VLAN.

Syntax Description

name	Specifies an ESRP-enabled or VRRP VLAN name.
vlan_tracked	Specifies the VLAN to be tracked.

Default

N / A.

Usage Guidelines

This command can be used with both ESRP-enabled VLANs and VRRP VLANs.

Example

The following command disables the tracking of port connectivity to VLAN *engineering*:

```
config vlan esrp-1 delete track-vlan engineering
```

History

This command was first available in ExtremeWare 4.1.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan esrp esrp-election

```
config vlan <name> esrp esrp-election [ports-track-priority |
ports-track-priority-mac | track-ports-priority | track-ports-priority-mac
| priority-ports-track-mac | priority-track-ports-mac | priority-mac-only]
```

Description

Configures the election algorithm on the switch.

Syntax Description

name	Specifies a VLAN name.
ports-track-priority	Specifies that this VLAN should consider election factors in the following order: Active ports, tracking information, ESRP priority.
ports-track-priority-mac	Specifies that this VLAN should consider election factors in the following order: Active ports, tracking information, ESRP priority, MAC address. This is the default election algorithm.
track-ports-priority	Specifies that this VLAN should consider election factors in the following order: Tracking information, active ports, ESRP priority.
track-ports-priority-mac	Specifies that this VLAN should consider election factors in the following order: Tracking information, active ports, ESRP priority, MAC address.
priority-ports-track-mac	Specifies that this VLAN should consider election factors in the following order: ESRP priority, active ports, tracking information, MAC address.
priority-track-ports-mac	Specifies that this VLAN should consider election factors in the following order: ESRP priority, tracking information, active ports, MAC address.
priority-mac-only	Specifies that this VLAN should consider election factors in the following order: ESRP priority, MAC address.

Default

`ports_track_priority_mac` election algorithm.

Usage Guidelines

The election algorithm determines the order of precedence of the election factors used to determine the ESRP Master. The election factors are:

- Active Ports (`ports`): the number of active ports (the switch with the highest number takes priority)
- Tracking Information (`track`): whether the switch is using ESRP tracking. A switch using tracking has priority.
- ESRP Priority (`priority`): a user-defined priority number between 0 and 254. A higher number has higher priority.
- MAC address (`mac`): the switch MAC address. A higher-number address has priority.

The election algorithm must be the same on all switches for a particular VLAN.

If no tracking information is configured for a particular field, the field is ignored.

The ports-track-priority or track-ports-priority options can be used to ensure that there is no failback if the original Master recovers (the Master will have the same ports, tracks and priority, but a higher MAC).

If a switch is master, it actively provides layer 3 routing services to other VLANs, and layer 2 switching between all the ports of that VLAN. Additionally, the switch exchanges ESRP packets with other switches that are in standby mode.

If a switch is in standby mode, it exchanges ESRP packets with other switches on that same VLAN. When a switch is in standby, it does not perform layer 3 routing or layer 2 switching services for the VLAN.

Example

The following command configures the election algorithm to use active port tracking information as the first criteria for determining the ESRP master switch for VLAN esrp-1:

```
config vlan esrp-1 esrp esrp-election track-ports-priority-mac
```

History

This command was first available in ExtremeWare 6.0.

The ports-track-priority and track-ports-priority election algorithms were added in ExtremeWare 6.2.1.

Platform Availability

This command is available on the "*i*" series platforms.

config vlan esrp priority

```
config vlan <name> esrp priority <value>
```

Description

Configures the ESRP priority.

Syntax Description

name	Specifies a VLAN name.
value	Specifies a number between 0 and 255.

Default

Priority = 0.

Usage Guidelines

The ESRP priority is one of the factors used by the ESRP election algorithm in determining which switch is the Master switch.

The range of the priority value is 0 to 254, with 0 being the lowest priority, 254 being the highest. If the ESRP priority is the determining criteria for the election algorithm, the highest priority value determines which switch will act as master for a particular VLAN.

Setting the priority to 255 configures the switch to standby mode, and to be ineligible to become the master. The switch will remain in standby mode even when the VLAN fails over from the current master. This feature is typically used to ensure a switch cannot become the ESRP master while it is offline for servicing.

Example

The following command configures the ESRP priority to the highest priority on VLAN `esrp-1`:

```
config vlan esrp-1 esrp priority 254
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

config vlan esrp timer

```
config vlan <name> esrp timer <timervalue> {esrp-nbr-timeout
<timeoutvalue>}
```

Description

Configures the ESRP timer values.

Syntax Description

name	Specifies a VLAN name.
timervalue	Specifies the number of seconds between keep-alive packets. The range is 1 to 255 seconds.
esrp-nbr-timeout	Specifies the number of seconds after which an ESRP neighbor times out. The range is 3 to 7650 seconds.

Default

The default `timervalue` is 2 seconds.

The default neighbor timeout is 3 times the `timervalue`.

Usage Guidelines

The timer specifies the interval, in seconds, for exchanging keep-alive packets between the ESRP switches for this VLAN. A lower value specifies a more frequent exchange of keep-alive messages, resulting in the faster detection of a failover condition. The timer setting must be configured identically for the VLAN across all participating switches.

The neighbor timeout specifies the amount of time that ESRP waits before considering the neighbor down. The timeout value must be at least 3 times, but not more than 30 times the timervalue. Entering a value outside of that range generates an error message.

Example

The following command configures the ESRP timer to 60 seconds:

```
config vlan esrp-1 esrp timer 60
```

History

This command was first available in ExtremeWare 4.0.

This command was modified to include the `esrp-nbr-timeout` option in ExtremeWare 6.2.2.

Platform Availability

This command is available on all platforms.

config vlan esrp group

```
config vlan <name> esrp group <group_number>
```

Description

Configures the group number to be used for the ESRP VLAN.

Syntax Description

name	Specifies a VLAN name.
group_number	Specifies the ESRP group to which this VLAN should be added.

Default

The default group number is 0.

Usage Guidelines

Each group runs an instance of ESRP within the same VLAN or broadcast domain. A maximum of four ESRP groups can be defined within the same networked broadcast domain. In addition a maximum of four groups are allowed per physical device (per switch)

The most typical application for multiple ESRP groups is when two or more sets of ESRP switches are providing fast-failover protection within a common subnet for two or more groups of users. An additional use for ESRP groups is ESRP Host Attach; ESRP VLANs that share ESRP HA ports must be members of different ESRP groups.

Example

The following command configures VLAN esrp-1 to be a member of ESRP group 2:

```
config vlan esrp-1 esrp group 2
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the "*i*" series platforms.

disable esrp vlan

```
disable esrp vlan <name>
```

Description

Disables ESRP on a VLAN.

Syntax Description

name	Specifies an ESRP-enabled VLAN name.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables ESRP on the VLAN *accounting*:

```
disable esrp vlan accounting
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

enable esrp vlan

```
enable esrp vlan <name>
```

Description

Enables ESRP on a VLAN.

Syntax Description

name	Specifies a VLAN name.

Default

Disabled.

Usage Guidelines

EDP must be enabled on all ports participating in ESRP.

A maximum of 3000 VLANs can run ESRP simultaneously on a single switch. A maximum of four switches can participate in providing redundant layer 3 or layer 2 services to a single VLAN.

ESRP cannot be enabled on the VLAN *default*.

Extreme switches that are not running ESRP, but are connected on a network that has other Extreme switches running ESRP are ESRP-*aware*. 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.

Example

The following command enables ESRP on the VLAN *esrp-1*:

```
enable esrp vlan esrp-1
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

show esrp

```
show esrp {detail}
```

Description

Displays ESRP configuration information.

Syntax Description

detail	Specifies detail for each switch in the ESRP VLAN.

Default

Shows summary ESRP information.

Usage Guidelines

This command shows information about the state of an ESRP VLAN and its neighbors. This includes:

- information about tracked devices

Example

The following command displays summary ESRP status information for the VLANs on the switch:

```
show esrp
```

It produces output similar to the following:

```
VLAN Name  VID Virtual IP/IPX   State Master MAC Addres Nbr Pri/Gr/Prt/TR/TP/T
 demo_esrp 4093 192.168.1.1      Slave 00:01:30:33:28:00 1   000/00/001/00/00/02

Nbr - Number of Neighbors, Pri - Priority In Use, Gr - Group, Prt - Number of
Active Ports, TR - Tracked Rt/Ping, TP - Tracked Ports, T - Hello Time.
        Host (Direct-attach, *=no-count) Ports on System:
```

The following command displays detailed ESRP status information for the VLANs on the switch:

```
show esrp detail
```

It produces output similar to the following:

```
VLAN Interface: demo_esrp.    Virtual IP address=192.168.1.1
        Priority:               0 (Priority In Use: 0)
        Active Ports:           1
        Tracked Rt/Ping:        0
        Tracked Ports:          0
        Tracked Diag:           -
        Tracked Env:            -
        Tracked RIP:            -
        Tracked OSPF:           -
        Tracked BGP:            -
        Election Algorithm:     ports-track-priority-mac
        Group:                  0
        Hello Timer:            2
        State:                  Enabled(Slave) on Wed Jan 23 10:09:48 2002
```

```
State Trans Counters:   ToMaster:(0)    ToSlave:(1)
Host ports(*=no-count): None
Restart Ports: None
Tracked VLANs: None
Tracked Ip Routes: None
Tracked Pings/Freq/N_miss: None
Neighbours:
[1]     Nbr Active Ports:       1
        Nbr Tracked Rt/Ping:    0
        Nbr Tracked Ports:      0
        Nbr Priority:           0
        Nbr MacID:              00:01:30:33:28:00
        Nbr HelloTimer:         2
        Nbr ESRP State:         Master
Host (Direct-attach, *=no-count) Ports on System:
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

show esrp vlan

```
show esrp vlan <name> {counters}
```

Description

Displays ESRP configuration information for a specific VLAN.

Syntax Description

name	Specifies an ESRP-enabled VLAN name.
counters	Displays ESRP counters.

Default

Displays summary information for the VLAN.

Usage Guidelines

None.

Example

The following command displays ESRP status information for ESRP-enabled VLAN *demo-esrp*:

```
show esrp vlan demo-esrp
```

It produces output similar to the following:

```
VLAN Interface: demo_esrp.    Virtual IP address=192.168.1.1
        Priority:               0 (Priority In Use: 0)
        Active Ports:           1
        Tracked Rt/Ping:        0
        Tracked Ports:          0
        Tracked Diag:           -
        Tracked Env:            -
        Tracked RIP:            -
        Tracked OSPF:           -
        Tracked BGP:            -
        Election Algorithm:     ports-track-priority-mac
        Group:                  0
        Hello Timer:            2
        State:                  Enabled(Slave) on Wed Jan 23 10:09:48 2002
        State Trans Counters:   ToMaster:(0)    ToSlave:(1)
        Host ports(*=no-count): None
        Restart Ports: None
        Tracked VLANs: None
        Tracked Ip Routes: None
        Tracked Pings/Freq/N_miss: None
        Neighbours:
        [1]     Nbr Active Ports:       1
                Nbr Tracked Rt/Ping:    0
                Nbr Tracked Ports:      0
                Nbr Priority:           0
```

```
Nbr MacID:              00:01:30:33:28:00
Nbr HelloTimer:         2
Nbr ESRP State:         Master
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

15 VRRP Commands

This chapter describes the following commands:

- Commands for enabling and disabling Virtual Router Redundancy Protocol (VRRP)
- Commands for performing basic VRRP configuration

NOTE

Commands for enabling and disabling port restart and enabling and disabling failure tracking for VRRP are described in Chapter 14, covering ESRP commands.

Like ESRP, VRRP is a protocol that allows multiple switches to provide redundant routing services to users. A virtual router is a group of one or more physical devices that acts as the default gateway for hosts on the network. The virtual router is identified by a virtual router identifier (VRID) and an IP address. All of the VRRP routers that participate in the virtual router are assigned the same VRID.

Extreme Networks' VRRP implementation is compliant with RFC 2338, Virtual Router Redundancy Protocol.

The following points pertain to VRRP:

- VRRP packets are encapsulated IP packets.
- The VRRP multicast address is 224.0.0.18.
- The virtual router MAC address is `00 00 5E 00 01 <vrid>`
- An interconnect link between VRRP routers should not be used, except when VRRP routers have hosts directly attached.
- A maximum of 64 VRID instances are supported on the router.
- Up to 4 unique VRIDs can be configured on an interface. VRIDs can be re-used, but not on the same interface.
- VRRP and Spanning Tree can be simultaneously enabled on the same switch.
- VRRP and ESRP cannot be simultaneously enabled on the same switch.

VRRP uses an election algorithm to dynamically assign responsibility for the master router to one of the VRRP routers on the network. A VRRP router is elected master if one of the following is true:

- The router is the IP address owner (router that has the IP address of the virtual router configured as its real interface address).
- The router is configured with the highest priority (the range is 1 - 255).

If the master router becomes unavailable, the election process provides dynamic failover and the backup router that has the highest priority assumes the role of master.

A new master is elected when one of the following things happen:

- VRRP is disabled on the master router.
- Communication is lost between master and backup router(s). The master router sends periodic advertisements to the backup routers to indicate that it is alive.

VRRP also supports the following tracking options:

- VRRP VLAN tracking
- VRRP route table tracking
- VRRP ping tracking

If a tracking option is enabled, and the object being tracked becomes unreachable, the master device will fail over. These tracking features are documented in the chapter on ESRP.

VRRP also supports port restart. Like the tracking features, the commands to enable and disable this feature are described in the chapter on ESRP.

config vrrp add vlan

```
config vrrp add vlan <name>
```

Description

Enables VRRP on a particular VLAN.

Syntax Description

name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

None.

Example

The following enables VRRP on VLAN *vrrp-1*:

```
config vrrp add vlan vrrp-1
```

History

This command was first available in ExtremeWare 6.2

Platform Availability

This command is available on the "*i*" series platforms.

config vrrp delete

```
config vrrp delete [vlan <name> | all]
```

Description

Disables VRRP on one or all VLANs.

Syntax Description

name	Specifies the name of a VLAN on which to disable VRRP.
all	Specifies that VRRP should be disabled on all VLANs on this device.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables VRRP on VLAN *vrrp-1*:

```
config vrrp delete vlan vrrp-1
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config vrrp vlan add

```
config vrrp vlan <name> add [master | backup] vrid <number> <ip address>
```

Description

Configures the VRID instance on the VRRP VLAN as master or backup.

Syntax Description

name	Specifies the name of a VRRP VLAN.
master	Specifies that this device is the master router for the virtual router.
backup	Specifies that this device is a backup router for this VLAN.
number	Specifies a Virtual Router ID (VRID). Value can be in the range of 1-255.
ip address	Specifies the IP address of the virtual router in which this device participates.

Default

N / A.

Usage Guidelines

The IP address must be the same on all VRRP routers that make up the virtual router for this VLAN. If the IP address is the same as the actual interface address of the device, this device is the IP address owner, and is automatically elected as the master router as long as it remains functional.

Example

The following command sets up this device as the master router for VLAN *vrrp-1*, using IP address 192.168.1.3 as the virtual router IP address:

```
config vrrp vlan vrrp-1 add master vrid 1 192.168.1.3
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config vrrp vlan authentication

```
config vrrp vlan <name> authentication [none | simple-password <simple
password>]
```

Description

Configures VRRP authentication.

Syntax Description

name	Specifies the name of a VRRP VLAN.
none	Specifies that no password is required.
simple password	Specifies the password for VRRP authentication. The maximum password length is eight characters.

Default

N/A.

Usage Guidelines

None.

Example

The following command configures authentication for VRRP VLAN *vrrp-1* with the password newvrrp:

```
config vrrp vlan vrrp-1 authentication simple-password newvrrp
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config vrrp vlan delete vrid

```
config vrrp vlan <name> delete vrid [<number> | all]
```

Description

Deletes one or all VRIDs.

Syntax Description

name	Specifies the name of a VRRP VLAN.
number	Specifies a Virtual Router ID (VRID). Value can be in the range of 1-255.
all	Specifies that all virtual routers should be deleted for this VLAN on this device.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the virtual router identified by VRID 2:

```
config vrrp vlan vrrp-1 delete vrid 2
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

config vrrp vlan vrid

```
config vrrp vlan <name> vrid <number> [priority <priority_number> |
advertisement-interval <ad_interval_number> | dont_preempt | preempt]
```

Description

Configures VRRP parameters.

Syntax Description

name	Specifies the name of a VRRP VLAN.
number	Specifies a Virtual Router ID (VRID). Value can be in the range of 1-255.
priority_number	Specifies the priority value to be used by this VRRP router in the master election process. The range is 1 - 254. The default value is 100.
ad_interval_number	Specifies the time interval between advertisements, in seconds. The range is 1 - 255. The default value is 1 second.
dont_preempt	Specifies that this router, as master, may not be preempted by a higher priority backup router.
preempt	Specifies that this router, as master, may be preempted by a higher-priority backup router. This is the default.

Default

N/A.

Usage Guidelines

This command may be used to configure a VRRP router priority, advertisement interval, and preempt mode.

The priority is used to determine which VRRP router takes over when the master fails over. A value of 255 is reserved for the router that is configured with the virtual router IP address. A value of 0 is reserved for the master router's use to indicate it is releasing responsibility for the virtual router.

The advertisement interval specifies the interval between advertisements sent by the master router to inform the backup routers that it is alive. The master down interval is the interval that a backup router waits after the last received advertisement before it determines that the master router is down.

The preempt mode controls whether a higher priority backup router preempts a lower priority master. `preempt` allows preemption. `dont_preempt` prohibits preemption. The default setting is `preempt`. The router that owns the virtual router IP address always preempts, independent of the setting of this parameter.

Example

The following commands set a priority and advertisement interval for the VRRP router on VLAN vrrp-1, and sets the preempt mode to disallow preemption:

```
config vrrp vlan vrrp-1 vrid 2 priority 200
config vrrp vlan vrrp-1 vrid 2 advertisement-interval 15
config vrrp vlan vrrp-1 vrid 2 dont_preempt
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

disable vrrp

```
disable vrrp
```

Description

Disables VRRP on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This disables VRRP on the device. All virtual routers defined on this device will also be disabled.

Example

The following command disables VRRP on the device:

```
disable vrrp
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

enable vrrp

```
enable vrrp
```

Description

Enables VRRP on the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command enables VRRP on this device:

```
enable vrrp
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

show vrrp

```
show vrrp [vlan <name> | all] {detail}
```

Description

Displays VRRP configuration information for one or all VLANs.

Syntax Description

name	Specifies the name of a VRRP VLAN.
all	Specifies that information should be displayed for all VLANs.
detail	Specifies detail information.

Default

N/A.

Usage Guidelines

Use the `detail` option for a detailed display.

Example

The following command displays summary status information for VRRP:

```
show vrrp
```

It produces output similar to the following:

```
VRRP Router: Enabled
   VLAN Name VRID Pri Virtual IP Addr State Master Mac Addres Prt/TR/TPr/W/M/T
 demo_vr(En) 0001 100 192.168.1.1      MSTR 00:00:5E:00:01:01   1  0   0 Y Y 1
```

En-Enabled, Ds-Disabled, Pri-Priority, T-Advert Timer, M-Preempt

Prt-Active Ports, TR-Tracked Routes/Pings, TPr-Tracked Ports, W-TrackWinner

The following command displays detail status information for VRRP:

```
show vrrp detail
```

It produces output similar to the following:

```
VRRP Router: Enabled
 Vlan:demo_vrrp IpAddress Owner=192.168.1.2 Vrrp:ENABLED Router:ENABLED
 Authentication: None
 Tracked VLANs:         -
 Tracked Ip Routes:     -
 Tracked Pings/Freq/N_miss:  -
 Tracked Diag:          -
 Tracked Env:           -
 Track Winner:          Yes
    1) Backup-Vrid:1 Virtual-IP:192.168.1.1 Priority:100
       Active Ports:1, Advert-Interval:1, Preempt:Yes
       State:MASTER on Wed Jan 23 10:17:42 2002
```

```
Transition Counters: ToMaster:1  ToBackup:1
Skew:0.609375 Master-Dn-Int:3.60938
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

show vrrp vlan stats

```
show vrrp vlan <name> stats
```

Description

Displays VRRP statistics for a particular VLAN.

Syntax Description

name	Specifies the name of a VRRP VLAN.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays statistics for VLAN *vrrp-1*:

```
show vrrp vlan vrrp-1 stats
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on the "*i*" series platforms.

16 IP Unicast Commands

Extreme Networks switches provide full layer 3, IP unicast routing. They exchange routing information with other routers on the network using either the Routing Information Protocol (RIP) or the Open Shortest Path First (OSPF) protocol. The switches dynamically build and maintain routing tables and determine the best path for each of its routes.

Each host that uses 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.

The routing software and hardware directs 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. The VLAN switching and IP routing functions occur within the switch.

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.

The Extreme Networks switch maintains an IP routing table for 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

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 are manually entered into the routing table. Static routes are used to reach networks not advertised by routers. You can configure up to 64 static unicast routes on the switch.

Static routes can also be used for security reasons, to control which routes you want advertised by the router. 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.

When there are multiple, conflicting choices of a route to a particular destination, the router picks the route with the longest matching network mask. If these are still equal, the router picks the route using the following criteria (in the order specified):

- Directly attached network interfaces
- ICMP redirects
- Static routes
- Directly attached network interfaces that are not active

If you define multiple default routes, the route that has the lowest metric is used. If there are multiple default routes that 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.

Internet Control Message Protocol (ICMP) is used to transmit information needed to control IP traffic. It is used mainly to provide information about routes to destination addresses. ICMP redirect messages inform hosts about more accurate routes to other systems, whereas ICMP unreachable messages indicate problems with a route.

Additionally, ICMP can cause TCP connection to terminate gracefully if the route becomes unavailable.

After IP unicast routing has been configured, you can configure the switch to forward Dynamic Host Configuration Protocol (DHCP) or BOOTP requests coming from clients on subnets being service 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.

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, the packet 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 can continue to use them.

To configure UPD-forwarding, you must first 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 (STD).

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 10 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

use a single UDP-forwarding profile. UDP packets directed toward a VLAN use an all-ones broadcast on that VLAN.

Proxy Address Resolution Protocol (ARP) was first developed 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 Extreme Networks switch supports proxy ARP for this type of network configuration.

Once IP ARP is 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).

After all the proxy ARP conditions have been met, the switch formulates an ARP Response using the configured MAC address in the packet.

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.

clear iparp

```
clear iparp {<ip address> | vlan <name>}
```

Description

Removes dynamic entries in the IP ARP table.

Syntax Description

ip address	Specifies an IP address.
name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

Permanent IP ARP entries are not affected.

Example

The following command removes a dynamically created entry from the IPARP table:

```
clear iparp 10.1.1.5/24
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

clear ipfdb

```
clear ipfdb {<ip address> <netmask>| vlan <name>}
```

Description

Removes the dynamic entries in the IP forwarding database.

Syntax Description

ip address	Specifies an IP address.
netmask	Specifies a subnet mask.
name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

If no options are specified, all IP FDB entries are removed.

Example

The following command removes dynamically created entries in the IP forwarding database:

```
clear ipfdb 10.1.2.1/24
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config bootprelay add

```
config bootprelay add <ip address>
```

Description

Configures the addresses to which BOOTP requests should be directed.

Syntax Description

ip address	Specifies an IP address.

Default

N/A.

Usage Guidelines

After IP unicast routing has been 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. 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 <ip address>
   ```

Example

The following command configures BOOTP requests to be directed to 123.45.67.8:

```
config bootprelay add 123.45.67.8
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms

config bootprelay delete

```
config bootprelay delete [<ip address> | all]
```

Description

Removes one or all IP destination addresses for forwarding BOOTP packets.

Syntax Description

ip address	Specifies an IP address.
all	Specifies all IP address entries.

Default

N/A.

Usage Guidelines

After IP unicast routing has been 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. 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 <ip address>
   ```

Example

The following command removes the destination address:

```
config bootprelay delete 123.45.67.8
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config iparp add

```
config iparp add <ip address> <mac_address>
```

Description

Adds a permanent entry to the ARP table. Specify the IP address and MAC address of the entry.

Syntax Description

ip address	Specifies an IP address.
mac_address	Specifies a MAC address.

Default

N / A.

Usage Guidelines

Add a permanent IP ARP entry to the system. The `ip address` is used to match the IP interface address to locate a suitable interface.

Example

The following command adds a permanent IP ARP entry to the switch for IP address *10.1.2.5*:

```
config iparp add 10.1.2.5 00:11:22:33:44:55
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config iparp add proxy

```
config iparp add proxy <ip address> {<mask>} {<mac_address>} {always}
```

Description

Configures the switch to respond to ARP Requests on behalf of devices that are incapable of doing so. Up to 64 proxy ARP entries can be configured.

Syntax Description

ip address	Specifies an IP address.
mask	Specifies a subnet mask.
mac_address	Specifies a MAC address.
always	Specifies all ARP Requests.

Default

N/A.

Usage Guidelines

When `mask` is not specified, an address with the mask 255.255.255.255 is assumed. When `mac_address` is not specified, the MAC address of the switch is used in the ARP Response. When `always` is specified, the switch answers ARP Requests without filtering requests that belong to the same subnet of the receiving router interface.

After IP ARP is 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).

After all the proxy ARP conditions have been met, the switch formulates an ARP Response using the configured MAC address in the packet.

Example

The following command configures the switch to answer ARP Requests for all devices with the address range of 100.101.45.1 to 100.101.45.255:

```
config iparp add proxy 100.101.45.0/24
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config iparp delete

```
config iparp delete <ip address>
```

Description

Deletes an entry from the ARP table. Specify the IP address of the entry.

Syntax Description

ip address	Specifies an IP address.

Default

N/A.

Usage Guidelines

Removes any IP ARP entry (dynamic or permanent) from the table. The `ip address` is used to match the IP interface address to locate a suitable interface.

Example

The following command deletes an IP address entry from the ARP table:

```
config iparp delete 10.1.2.5
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config iparp delete proxy

```
config iparp delete proxy [<ip address> {<mask>} | all]
```

Description

Deletes one or all proxy ARP entries.

Syntax Description

ip address	Specifies an IP address.
mask	Specifies a subnet mask.
all	Specifies all ARP entries.

Default

Not Always.

Usage Guidelines

Proxy ARP can be used for two purposes:

1. To support host that cannot process ARP traffic. In this case, the switch answers the ARP Request for that host.
2. To hide the IP topology from the host. The network administrator can configure a large network on the host machine (16-bit mask) and a smaller network on each router interface (for example, 22-bit mask). When the host sends ARP Request for another host on another subnet, the switch answers the ARP Request and all subsequent traffic will be sent directly to the router.

You can configure up to 64 proxy ARP entries. When the `mask` is not specified, then software will assume a host address (that is, a 32-bit mask). When the MAC address is not specified, then the software uses the switch's MAC address as the proxy host. Always should be specified for type-1 usage, not always is the default (type-2).

Example

The following command deletes the IP ARP proxy entry 1*00.101.45.0/24*:

```
config iparp delete proxy 100.101.45.0/24
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config iparp max-entries

```
config iparp max-entries <number>
```

Description

Configures the maximum allowed IP ARP entries.

Syntax Description

number	Specifies a number of maximum IP ARP entries.

Default

4096.

Usage Guidelines

Range: 1 - 20480. The maximum IP ARP entries include dynamic, static, and incomplete IP ARP entries.

Example

The following command sets the maximum IP ARP entries to 2000 entries:

```
config iparp max-entries 2000
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all platforms.

config iparp max-pending-entries

```
config iparp max-pending-entries <number>
```

Description

Configures the maximum allowed incomplete IP ARP entries.

Syntax Description

number	Specifies a number of maximum IP ARP entries.

Default

256.

Usage Guidelines

Range: 1 - 20480, but cannot be greater than the configured IP ARP max-entries value.

Example

The following command sets the maximum IP ARP entries to 500 entries:

```
config iparp max-pending-entries 500
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all platforms.

config iparp timeout

```
config iparp timeout <minutes>
```

Description

Configures the IP ARP timeout period.

Syntax Description

minutes	Specifies a time in minutes.

Default

20 minutes.

Usage Guidelines

A setting of 0 disables ARP aging.

Example

The following command sets the IP ARP timeout period to 10 minutes:

```
config iparp timeout 10
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

config ip-down-vlan-action

```
config ip-down-vlan-action [consume | drop | forward]
```

Description

Configures the forwarding functionality destined to nonworking IP interfaces.

Syntax Description

consume	Specifies the consume function.
drop	Specifies the drop function.
forward	Specifies the forwarding function.

Default

N/A.

Usage Guidelines

None.

Example

The following command configures the forwarding functionality destined to nonworking IP interfaces:

```
config ip-down-vlan-action forward
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all platforms.

config iproute add

```
config iproute add <ip address> <mask> <gateway> <metric>
```

Description

Adds a static address to the routing table.

Syntax Description

ip address	Specifies an IP address.
mask	Specifies a subnet mask.
gateway	Specifies a VLAN gateway.
metric	Specifies a cost metric.

Default

N/A.

Usage Guidelines

Use a value of 255.255.255.255 for mask to indicate a host entry.

Example

The following command adds a static address to the routing table:

```
config iproute add 10.1.1.1/24 123.45.67.1 5
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config iproute add blackhole

```
config iproute add blackhole <ip address> <mask>
```

Description

Adds a blackhole address to the routing table. All traffic destined for a configured blackhole IP address is silently dropped, and no Internet Control Message Protocol (ICMP) message is generated.

Syntax Description

ip address	Specifies an IP address.
mask	Specifies a subnet mask.

Default

N/A.

Usage Guidelines

A blackhole entry configures packets with a specified MAC destination address to be discarded. 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 forwarding database (FDB).

Example

The following command adds a blackhole address to the routing table for packets with a destination address of 100.101.145.4:

```
config iproute add blackhole 100.101.145.4
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config iproute add blackhole default

```
config iproute add blackhole default
```

Description

Adds a default blackhole route to the routing table. All traffic destined for an unknown IP destination is silently dropped, and no Internet Control Message Protocol (ICMP) message is generated.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

While a default route is for *forwarding* traffic destined to an unknown IP destination, and a blackhole route is for *discarding* traffic destined to a specified IP destination, a *default blackhole* route is for *discarding* traffic to the unknown IP destination.

Using this command, all traffic with an unknown destination is discarded. If there is another static default route existing in the routing table, the `blackhole default` route takes higher route priority.

The default blackhole route is treated like a permanent entry in the event of a switch reset or power off/on cycle. The default blackhole route's origin is "b" or "blackhole" and the gateway IP address for this route is 0.0.0.0.

Example

The following command adds a blackhole default route into the routing table:

```
config iproute add blackhole default
```

History

This command was first available in ExtremeWare 6.2.2.

Platform Availability

This command is available on all platforms.

config iproute add default

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

Description

Adds a default gateway to the routing table.

Syntax Description

gateway	Specifies a VLAN gateway
metric	Specifies a cost metric. If no metric is specified, the default of 1 is used.

Default

If no metric is specified, the default metric of 1 is used.

Usage Guidelines

Default routes are used when the router has no other dynamic or static route to the requested destination. A default gateway must be located on a configured IP interface. 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.

Example

The following command configures a default route for the switch:

```
config iproute add default 123.45.67.1
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config iproute delete

```
config iproute delete <ip address> <mask> <gateway>
```

Description

Deletes a static address from the routing table.

Syntax Description

ip address	Specifies an IP address.
mask	Specifies a subnet mask.
gateway	Specifies a VLAN gateway.

Default

N / A.

Usage Guidelines

Use a value of 255.255.255.255 for mask to indicate a host entry.

Example

The following command deletes an address from the gateway:

```
config iproute delete 10.101.0.250/24 10.101.0.1
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config iproute delete blackhole

```
config iproute delete blackhole <ip address> <mask>
```

Description

Deletes a blackhole address from the routing table.

Syntax Description

ip address	Specifies an IP address.
mask	Specifies a subnet mask.

Default

N/A.

Usage Guidelines

None.

Example

The following command removes a blackhole address from the routing table:

```
config iproute delete blackhole 100.101.145.4
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config iproute delete blackhole default

```
config iproute delete blackhole default
```

Description

Deletes a default blackhole route from the routing table.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes a blackhole default route from the routing table:

```
config iproute delete blackhole default
```

History

This command was first available in ExtremeWare 6.2.2.

Platform Availability

This command is available on all platforms.

config iproute delete default

```
config iproute delete default <gateway>
```

Description

Deletes a default gateway from the routing table.

Syntax Description

gateway	Specifies a VLAN gateway.

Default

N / A.

Usage Guidelines

Default routes are used when the router has no other dynamic or static route to the requested destination. A default gateway must be located on a configured IP interface.

Example

The following command deletes a default gateway:

```
config iproute delete default 123.45.67.1
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config iproute priority

```
config iproute priority [rip | bootp | icmp | static | ospf-intra |
ospf-inter | ospf-as-external | ospf-extern1 | ospf-extern2] <priority>
```

Description

Changes the priority for all routes from a particular route origin.

Syntax Description

rip	Specifies RIP.
bootp	Specifies BOOTP.
icmp	Specifies ICMP.
static	Specifies static routes.
ospf-intra	Specifies OSPFIntra routing.
ospf-inter	Specifies OSPFInter routing.
ospf-as-external	Specifies OSPF as External routing.
ospf-extern1	Specifies OSPF External 1 routing.
ospf-extern2	Specifies OSPF External 2 routing.
priority	Specifies a priority number.

Default

Table 16 lists the relative priorities assigned to routes depending upon the learned source of the route.

Table 16: Relative Route Priorities

Route Origin	Priority
Direct	10
Blackhole	50
Static	1100
ICMP	1200
OSPFIntra	2200
OSPFInter	2300
RIP	2400
OSPF External 1	3200
OSPF External 2	3300
BOOTP	5000

Usage Guidelines

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

Example

The following command sets IP route priority for static routing to 1200:

```
config iproute priority static 1200
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

config iproute route-map

```
config iproute route-map [bgp | direct | e-bgp | i-bgp | ospf |
ospf-extern1 | ospf-extern2 | ospf-inter | ospf-intra | rip | static]
[<route map> | none]
```

Description

Configures the contents of the IP routing table.

Syntax Description

bgp	Specifies BGP routing.
direct	Specifies direct routing.
e-bgp	Specifies E-BGP routing.
i-bgp	Specifies I-BGP routing.
opsf	Specifies OSPF routing.
ospf-extern1	Specifies OSPF External 1 routing.
ospf-extern2	Specifies OSPF External 2 routing.
ospf-inter	Specifies OSPFInter routing.
ospf-intra	Specifies OSPFIntra routing.
rip	Specifies RIP routing.
static	Specifies static routing.
route map	Specifies a route map.
none	Specifies not to use a route map.

Default

N / A.

Usage Guidelines

Route maps for IP routing can be configured based on the route origin. When routes are added to the IP routing table from various sources, the route map configured for the origin of the route is applied to the route. After matching on specified characteristics, the characteristics for the route can be modified using the route maps. The characteristics that can be matched and modified are dependent on the origin of the route. Route maps for IP routing can be dynamically changed. In the case of direct and static route origins, the changes are reflected immediately. In the case of routes that are sourced from other origin, the changes are reflected within 30 seconds.

MPLS uses route map-based filters for controlling label advertisement and label propagation. The implementation of the `delete route-map <route-map>` command has been augmented to support the MPLS module.

Example

The following command configures the IP routing table *bgp_out* to BGP routing:

```
config iproute route-map bgp_out bgp
```

History

This command was first available in ExtremeWare 6.1.5.

This command was modified in an ExtremeWare IP Technology Services Release based on ExtremeWare v6.1.8b12 to support MPLS modules.

Platform Availability

This command is available on "*i*" series platforms.

config irdp

```
config irdp [multicast | broadcast]
```

Description

Configures the destination address of the router advertisement messages.

Syntax Description

multicast	Specifies multicast setting.
broadcast	Specifies broadcast setting.

Default

Multicast (224.0.0.1).

Usage Guidelines

None.

Example

The following command sets the address of the router advertiser messages to multicast:

```
config irdp multicast
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

config irdp

```
config irdp <mininterval> <maxinterval> <lifetime> <preference>
```

Description

Configures the router advertisement message timers, using seconds.

Syntax Description

mininterval	Specifies the minimum amount of time between router advertisements in seconds. The default setting is 450 seconds.
maxinterval	Specifies the maximum amount of time between router advertisements in seconds. The default setting is 600 seconds.
lifetime	Specifies the client aging time. The default setting is 1,800 seconds.
preference	Specifies the preference level of the router. The default setting is 0.

Default

N/A.

Usage Guidelines

All arguments need to be specified. All time intervals are in seconds.

An ICMP Router Discover Protocol (IRDP) client always uses the router with the highest preference level. Change the preference setting to encourage or discourage the use of this router. The default setting is 0.

Example

The following command configures the router advertisement message timers:

```
config irdp 30 40 300 1
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

config tcp-sync-rate

```
config tcp-sync-rate <number_sync_per_sec>
```

Description

Configures a limit for the switch to process TCP connection requests.

Syntax Description

number_sync_per_sec	Specifies a time in seconds.

Default

25.

Usage Guidelines

If the connection request rate is higher than the specified rater, or the total number of outstanding connection requests exceed the system limit, the system ages out incomplete connection requests at a faster rate. The range is from 5 to 200,000.

Example

The following command configures a 50 second limit for the switch to process TCP connection requests:

```
config tcp-sync-rate 50
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on "*i*" series platforms.

config udp-profile add

```
config udp-profile <profile_name> add <udp_port> [vlan <name> | ip address
<dest_ipaddress>]
```

Description

Configures a UDP-forwarding profile.

Syntax Description

profile_name	Specifies a UDP profile name.
udp_port	Specifies a UDP port number.
name	Specifies a VLAN name.
dest_ipaddress	Specifies an IP address.

Default

N / A.

Usage Guidelines

A maximum of 10 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.

Example

The following command adds port 34 to UDP profile *port_34_to_server*:

```
config udp-profile port_34_to_server add 34 ip address 10.1.1.1
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on “*i*” series platforms.

config udp-profile delete

```
config udp-profile <profile_name> delete <udp_port> [vlan <name> | ip
address <dest_ipaddress>]
```

Description

Deletes a forwarding entry from the specified UDP-profile.

Syntax Description

profile_name	Specifies a UDP profile name.
udp_port	Specifies a UDP port number.
name	Specifies a VLAN name.
dest_ipaddress	Specifies an IP address.

Default

N/A.

Usage Guidelines

None.

Example

The following command removes port 34 from UDP profile *port_34_to_server*:

```
config udp-profile port_34_to_server delete 34 ip address 10.1.1.1
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

config vlan subvlan address range

```
config vlan <name> subvlan-address-range <ip address1> - <ip address2>
```

Description

Configures sub-VLAN address ranges on each sub-VLAN to prohibit the entry of IP addresses from hosts outside of the configured range.

Syntax Description

name	Specifies a super-VLAN name.
ip address1	Specifies an IP address.
ip address2	Specifies another IP address.

Default

N / A.

Usage Guidelines

There is no error checking to prevent the configuration of overlapping sub-VLAN address ranges between multiple sub-VLANs. Doing so can result in unexpected behavior of ARP within the super-VLAN and associated sub-VLANs.

Example

The following command configures the super-VLAN *vsuper* to prohibit the entry of IP addresses from hosts outside of the configured range of IP addresses:

```
config vlan vsuper subvlan-address-range 10.1.1.1 - 10.1.1.255
```

History

This command was available in ExtremeWare 6.1.5.

Platform Availability

This command is available on "*i*" series platforms.

config vlan upd-profile

```
config vlan <name> udp-profile <profile_name>
```

Description

Assigns a UDP-forwarding profile to the source VLAN.

Syntax Description

name	Specifies a VLAN name.
profile_name	Specifies a UDP profile name.

Default

N/A.

Usage Guidelines

After the UDP profile has been associated with the VLAN, the switch picks up any broadcast UDP packets that match 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 BOOTP/DHCP proxy functions are invoked.

Example

The following command assigns a UDP profile to VLAN *accounting*:

```
config vlan accounting udp-profile port_34_to_server
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on "*i*" series platforms.

config vlan secondary-ip

```
config vlan <super-vlan name> [add | delete] secondary-ip <ip address>
{<mask>}
```

Description

Adds or deletes a secondary IP address to the super-VLAN for responding to ICMP ping requests.

Syntax Description

super-vlan name	Specifies a super-VLAN name.
add	Specifies to add a secondary IP address.
delete	Specifies to delete a secondary IP address.
ip address	Specifies an IP address.
mask	Specifies a netmask.

Default

N/A.

Usage Guidelines

- All broadcast and unknown traffic remains local to the sub-VLAN and does not cross the sub-VLAN boundary. All traffic within the sub-VLAN is switched by the sub-VLAN, allowing traffic separation between sub-VLANs (while using the same default router address among the sub-VLANs).
- Hosts are located on the sub-VLAN. Each host can assume any IP address within the address range of the super-VLAN router interface. Hosts on the sub-VLAN are expected to have the same network mask as the super-VLAN and have their default router set to the IP address or the super-VLAN.
- All traffic (IP unicast and IP multicast) between sub-VLANs is routed through the super-VLAN. For example, no ICMP redirects are generated for traffic between sub-VLANs, because the super-VLAN is responsible for sub-VLAN routing. Unicast IP traffic across the sub-VLANs is facilitated by the automatic addition of an ARP entry (similar to a proxy ARP entry) when a sub-VLAN is added to a super-VLAN. This feature can be disabled for security purposes.

IP multicast traffic between sub-VLANs is routed when an IP multicast routing protocol is enabled on the super-VLAN.

Example

The following command adds a secondary IP address to the super-VLAN *vsuper* for responding to ICMP ping requests:

```
config vlan vsuper add secondary-ip 10.1.1.1
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

config vlan subvlan

```
config vlan <super-vlan name> [add | delete] subvlan <sub-vlan name>
```

Description

Adds or deletes a sub-VLAN to a super-VLAN.

Syntax Description

super-vlan name	Specifies a super-VLAN name
add	Specifies to add the sub-VLAN to the super-VLAN
delete	Specifies to delete the sub-VLAN from the super-VLAN
sub-vlan name	Specifies a sub-VLAN name.

Default

N / A.

Usage Guidelines

- All broadcast and unknown traffic remains local to the sub-VLAN and does not cross the sub-VLAN boundary. All traffic within the sub-VLAN is switched by the sub-VLAN, allowing traffic separation between sub-VLANs (while using the same default router address among the sub-VLANs).
- Hosts are located on the sub-VLAN. Each host can assume any IP address within the address range of the super-VLAN router interface. Hosts on the sub-VLAN are expected to have the same network mask as the super-VLAN and have their default router set to the IP address or the super-VLAN.
- All traffic (IP unicast and IP multicast) between sub-VLANs is routed through the super-VLAN. For example, no ICMP redirects are generated for traffic between sub-VLANs, because the super-VLAN is responsible for sub-VLAN routing. Unicast IP traffic across the sub-VLANs is facilitated by the automatic addition of an ARP entry (similar to a proxy ARP entry) when a sub-VLAN is added to a super-VLAN. This feature can be disabled for security purposes.

IP multicast traffic between sub-VLANs is routed when an IP multicast routing protocol is enabled on the super-VLAN.

Example

The following command adds the sub-VLAN *vsub1*to the super-VLAN *vsuper*:

```
config vlan vsuper add subvlan vsub1
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on the "*i*" series platforms.

create udp-profile

```
create udp-profile <profile_name>
```

Description

Creates a UDP-forwarding destination profile that describes the types of UDP packets (by port number) that are used, and where they are to be forwarded.

Syntax Description

profile_name	Specifies a UDP profile name.

Default

N/A.

Usage Guidelines

You must give the profile a unique name, in the same manner as a VLAN, protocol filter, or Spanning Tree Domain (STD). A maximum of 10 UDP-forwarding profiles can be defined.

Example

The following command creates a UPD profile named *backbone*:

```
create udp-profile backbone
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

delete udp-profile

```
delete udp-profile <profile_name>
```

Description

Deletes a UDP-forwarding profile.

Syntax Description

profile_name	Specifies a UDP profile name.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes a UPD profile named *backbone*:

```
delete udp-profile backbone
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

disable bootp vlan

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

Description

Disables the generation and processing of BOOTP packets on a VLAN to obtain an IP address for the VLAN from a BOOTP server.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

Enabled for all VLANs.

Usage Guidelines

None.

Example

The following command disables the generation and processing of BOOTP packets on a VLAN named *accounting*:

```
disable bootp vlan accounting
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

disable bootprelay

```
disable bootprelay
```

Description

Disables the BOOTP relay function.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

After IP unicast routing has been 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 <ip address>
   ```

Example

The following command disables the forwarding of BOOTP requests:

```
disable bootprelay
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

disable icmp address-mask

```
disable icmp address-mask {vlan <name>}
```

Description

Disables the generation of an ICMP address-mask reply on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

Disables 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.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command disables the generation of an ICMP address-mask reply on VLAN *accounting*:

```
disable icmp address-mask vlan accounting
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

disable icmp parameter-problem

```
disable icmp parameter-problem {vlan <name>}
```

Description

Disables the generation of an ICMP parameter-problem message on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

Disables the generation of an ICMP parameter-problem message (type 12) when the switch cannot properly process the IP header or IP option information. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command disables the generation of an ICMP parameter-problem message on VLAN *accounting*:

```
disable icmp parameter-problem vlan accounting
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

disable icmp port-unreachables

```
disable icmp port-unreachables {vlan <name>}
```

Description

Disables the generation of ICMP port unreachable messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

Disables the generation of ICMP port unreachable messages (type 3, code 3) when a TPC or UDP request is made to the switch, and no application is waiting for the request, or access policy denies the request. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command disables ICMP port unreachable messages on VLAN *accounting*:

```
disable icmp port-unreachables vlan accounting
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

disable icmp redirects

```
disable icmp redirects {vlan <name>}
```

Description

Disables generation of ICMP redirect messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

This option only applies to the switch when the switch is not in routing mode.

Example

The following command disables ICMP redirects from VLAN *accounting*:

```
disable icmp redirects vlan accounting
```

History

This command was available in ExtremeWare 2.0.

This command was modified in an ExtremeWare IP Technology Services Release based on ExtremeWare v6.1.8b12 to support MPLS modules.

Platform Availability

This command is available on all platforms.

disable icmp time-exceeded

```
disable icmp time-exceeded {vlan <name>}
```

Description

Disables the generation of ICMP time exceeded messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

Disables 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. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command disables the generation of ICMP time exceeded messages on VLAN *accounting*:

```
disable icmp time-exceeded vlan accounting
```

History

This command was first available in ExtremeWare 6.1.

This command was modified in an ExtremeWare IP Technology Services Release based on ExtremeWare v6.1.8b12 to support MPLS modules.

Platform Availability

This command is available on "*i*" series platforms.

disable icmp timestamp

```
disable icmp timestamp {vlan <name>}
```

Description

Disables the generation of an ICMP timestamp response on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

Disables the generation of an ICMP timestamp response (type 14, code 0) when an ICMP timestamp request is received. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command disables the generation of an ICMP timestamp response on VLAN *accounting*:

```
disable icmp timestamp vlan accounting
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

disable icmp unreachables

```
disable icmp unreachables {vlan <name>}
```

Description

Disables the generation of ICMP unreachable messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

None.

Example

The following command disables the generation of ICMP unreachable messages on all VLANs:

```
disable icmp unreachables
```

History

This command was available in ExtremeWare 2.0.

This command was modified in an ExtremeWare IP Technology Services Release based on ExtremeWare v6.1.8b12 to support MPLS modules.

Platform Availability

This command is available on "*i*" series platforms.

disable icmp useredirects

```
disable icmp useredirects
```

Description

Disables the modification of route table information when an ICMP redirect message is received.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This option only applies to the switch when the switch is not in routing mode.

Example

The following command disables the changing of routing table information:

```
disable icmp useredirects
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

disable iparp checking

```
disable iparp checking
```

Description

Disable checking if the ARP Request source IP address is within the range of the local interface or VLAN domain.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

None.

Example

The following command disables IP ARP checking:

```
disable iparp checking
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on "*i*" series platforms.

disable iparp refresh

```
disable iparp refresh
```

Description

Disables IP ARP to refresh its IP ARP entries before timing out.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

IP ARP refresh can only be disabled if IP forwarding is disabled. The purpose of disabling ARP refresh is to reduce ARP traffic in a high node count layer 2 switching only environment.

Example

The following command disables IP ARP refresh:

```
disable iparp refresh
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on "*i*" series platforms.

disable ipforwarding

```
disable ipforwarding {[broadcast | fast-direct-broadcast |
ignore-broadcast]} {vlan <name>}
```

Description

Disables routing (or routing of broadcasts) for one or all VLANs. If no argument is provided, disables routing for all VLANs.

Syntax Description

broadcast	Specifies broadcast IP forwarding.
fast-direct-broadcast	Specifies fast direct broadcast forwarding.
ignore-broadcast	Specifies to ignore broadcast forwarding.
name	Specifies a VLAN name.

Default

Disabled.

Usage Guidelines

Disabling IP forwarding also disables broadcast forwarding. Broadcast forwarding can be disabled without disabling IP forwarding. When new IP interfaces are added, IP forwarding (and IP broadcast forwarding) is disabled by default.

Other IP related configuration is not affected.

Example

The following command disables forwarding of IP broadcast traffic for a VLAN named *accounting*:

```
disable ipforwarding broadcast vlan accounting
```

History

This command was available in ExtremeWare 2.0.

This command was modified in an ExtremeWare IP Technology Services Release based on ExtremeWare v6.1.8b12 to support MPLS modules.

Platform Availability

This command is available on "*i*" series platforms.

disable ip-option loose-source-route

```
disable ip-option loose-source-route
```

Description

Disables the loose source route IP option.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

None.

Example

The following command disables the loose source route IP option:

```
disable ip-option loose-source-route
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on "*i*" series platforms.

disable ip-option record-route

```
disable ip-option record-route
```

Description

Disables the record route IP option.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

None.

Example

The following command disables the record route IP option:

```
disable ip-option record-route
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on "*i*" series platforms.

disable ip-option record-timestamp

```
disable ip-option record-timestamp
```

Description

Disables the record timestamp IP option.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

None.

Example

The following command disables the record timestamp IP option:

```
disable ip-option record-timestamp
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on "*i*" series platforms.

disable ip-option strict-source-route

```
disable ip-option strict-source-route
```

Description

Disables the strict source route IP option.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

None.

Example

The following command disables the strict source route IP option:

```
disable ip-option strict-source-route
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on "*i*" series platforms.

disable ip-option use-router-alert

```
disable ip-option use-router-alert
```

Description

Disables the generation of the router alert IP option.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

None.

Example

The following command disables generation of the router alert IP option:

```
disable ip-option use-router-alert
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on "*i*" series platforms.

disable iproute sharing

```
disable iproute sharing
```

Description

Disables load sharing if multiple routes to the same destination are available. When multiple routes to the same destination are available, load sharing can be enabled to distribute the traffic to multiple destination gateways. Only paths with the same lowest cost is will be shared.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

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.

Configure static routes and/or OSPF as you would normally. ExtremeWare supports unlimited route sharing across static routes and up to 12 ECMP routes for OSPF.

Route sharing is useful only in instances where you are constrained for bandwidth. This is typically not the case using Extreme switches. Using route sharing makes router troubleshooting more difficult because of the complexity in predicting the path over which the traffic will travel.

Example

The following command disables load sharing for multiple routes:

```
disable iproute sharing
```

History

This command was available in ExtremeWare 2.0.

This command was modified in ExtremeWare 6.2.2 to allow support of up to 12 ECMP routes for OSPF.

Platform Availability

This command is available on all platforms.

disable irdp

```
disable irdp {vlan <name>}
```

Description

Disables the generation of ICMP router advertisement messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

If no optional argument is specified, all the IP interfaces are affected.

Example

The following command disables IRDP on VLAN *accounting*:

```
disable irdp vlan accounting
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

disable loopback-mode vlan

```
disable loopback-mode vlan [<name> | all]
```

Description

Disallows a VLAN to be placed in the UP state without an external active port. This allows (disallows) the VLANs routing interface to become active.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

N / A.

Usage Guidelines

Use this command to specify a stable interface as a source interface for routing protocols. This decreases the possibility of route flapping, which can disrupt connectivity.

Example

The following command disallows the VLAN *accounting* to be placed in the UP state without an external active port:

```
disable loopback-mode vlan accounting
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

disable multinetting

```
disable multinetting
```

Description

Disables IP multinetting on the system.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The multinetting features requires the user to change the MAC FDB aging timer to be at least 3000 seconds on the switch. This command will automatically change the FDB timer to 3000 seconds if it is shorter than 3000 seconds.

Example

The following command disables multinetting on the system:

```
disable multinetting
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

disable subvlan-proxy-arp vlan

```
disable subvlan-proxy-arp vlan [<super-vlan name> | all]
```

Description

Disables the automatic entry of sub-VLAN information in the proxy ARP table.

Syntax Description

super-vlan name	Specifies a super-VLAN name.
all	Specifies all VLANs.

Default

Enabled.

Usage Guidelines

To facilitate communication between sub-VLANs, by default, an entry is made in the IP ARP table of the super-VLAN that performs a proxy ARP function. This allows clients on one sub-VLAN to communicate with clients on another sub-VLAN. In certain circumstances, intra-sub-VLAN communication may not be desired for isolation reasons.

NOTE

The isolation option works for normal, dynamic, ARP-based client communication.

Example

The following command disables the automatic entry of sub-VLAN information in the proxy ARP table of the super-VLAN *vsuper*:

```
disable subvlan-proxy-arp vlan vsuper
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

enable bootp vlan

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

Description

Enables the generation and processing of BOOTP packets on a VLAN to obtain an IP address for the VLAN from a BOOTP server.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

Enabled for all VLANs.

Usage Guidelines

None.

Example

The following command enables the generation and processing of BOOTP packets on a VLAN named *accounting*:

```
enable bootp vlan accounting
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

enable bootprelay

```
enable bootprelay
```

Description

Enables the BOOTP relay function.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

After IP unicast routing has been 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 <ip address>
   ```

Example

The following command enables the forwarding of BOOTP requests:

```
enable bootprelay
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

enable icmp address-mask

```
enable icmp address-mask {vlan <name>}
```

Description

Enables the generation of an ICMP address-mask reply on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

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.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command enables the generation of an ICMP address-mask reply on VLAN *accounting*:

```
enable icmp address-mask vlan accounting
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

enable icmp parameter-problem

```
enable icmp parameter-problem {vlan <name>}
```

Description

Enables the generation of an ICMP parameter-problem message on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

Enables the generation of an ICMP parameter-problem message (type 12) when the switch cannot properly process the IP header or IP option information. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command enables the generation of an ICMP parameter-problem message on VLAN *accounting*:

```
enable icmp parameter-problem vlan accounting
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

enable icmp port-unreachables

```
enable icmp port-unreachables {vlan <name>}
```

Description

Enables the generation of ICMP port unreachable messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

Enables the generation of ICMP port unreachable messages (type 3, code 3) when a TPC or UDP request is made to the switch, and no application is waiting for the request, or access policy denies the request. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command enables ICMP port unreachable messages on VLAN *accounting*:

```
enable icmp port-unreachables vlan accounting
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

enable icmp redirects

```
enable icmp redirects {vlan <name>}
```

Description

Enables generation of ICMP redirect messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

This option only applies to the switch when the switch is not in routing mode.

Example

The following command enables the generation of ICMP redirect messages on all VLANs:

```
enable icmp redirects
```

History

This command was available in ExtremeWare 2.0.

This command was modified in an ExtremeWare IP Technology Services Release based on ExtremeWare v6.1.8b12 to support MPLS modules.

Platform Availability

This command is available on all platforms.

enable icmp time-exceeded

```
enable icmp time-exceeded {vlan <name>}
```

Description

Enables the generation of ICMP time exceeded messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

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. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command enables the generation of ICMP time exceeded messages on VLAN *accounting*:

```
enable icmp time-exceeded vlan accounting
```

History

This command was first available in ExtremeWare 6.1.

This command was modified in an ExtremeWare IP Technology Services Release based on ExtremeWare v6.1.8b12 to support MPLS modules.

Platform Availability

This command is available on "*i*" series platforms.

enable icmp timestamp

```
enable icmp timestamp {vlan <name>}
```

Description

Enables the generation of an ICMP timestamp response on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

Enables the generation of an ICMP timestamp response (type 14, code 0) when an ICMP timestamp request is received. If a VLAN is not specified, the command applies to all IP interfaces.

This command only affects the generation of certain ICMP packets. Filtering of ICMP packets usually forwarded by the switch is controlled by the access-list commands.

Example

The following command enables the generation of an ICMP timestamp response on VLAN *accounting*:

```
enable icmp timestamp vlan accounting
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

enable icmp unreachables

```
enable icmp unreachables {vlan <name>}
```

Description

Enables the generation of ICMP unreachable messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

None.

Example

The following command enables the generation of ICMP unreachable messages on all VLANs:

```
enable icmp unreachables
```

History

This command was available in ExtremeWare 2.0.

This command was modified in an ExtremeWare IP Technology Services Release based on ExtremeWare v6.1.8b12 to support MPLS modules.

Platform Availability

This command is available on all platforms.

enable icmp useredirects

```
enable icmp useredirects
```

Description

Enables the modification of route table information when an ICMP redirect message is received.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This option only applies to the switch when the switch is not in routing mode.

Example

The following command enables the modification of route table information:

```
enable icmp useredirects
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

enable iparp checking

```
enable iparp checking
```

Description

Enables checking if the ARP Request source IP address is within the range of the local interface or VLAN domain.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

None.

Example

The following command enables IP ARP checking:

```
enable iparp checking
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on "*i*" series platforms.

enable iparp refresh

```
enable iparp refreshenable iparp refresh
```

Description

Enables IP ARP to refresh its IP ARP entries before timing out.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

IP ARP refresh can only be disabled if IP forwarding is disabled. The purpose of disabling ARP refresh is to reduce ARP traffic in a high node count layer 2 switching only environment.

Example

The following command enables IP ARP refresh:

```
enable iparp refresh
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on "*i*" series platforms.

enable ipforwarding

```
enable ipforwarding {[broadcast | fast-direct-broadcast |
ignore-broadcast]} {vlan <name>}
```

Description

Enables IP routing or IP broadcast forwarding for one or all VLANs. If no argument is provided, enables IP routing for all VLANs that have been configured with an IP address.

Syntax Description

broadcast	Specifies broadcast IP forwarding.
fast-direct-broadcast	Specifies fast direct broadcast forwarding.
ignore-broadcast	Specifies to ignore broadcast forwarding.
name	Specifies a VLAN name.

Default

Disabled.

Usage Guidelines

IP forwarding must first be enabled before IP broadcast forwarding can be enabled. When new IP interfaces are added, IP forwarding (and IP broadcast forwarding) is disabled by default.

Other IP related configuration is not affected.

Example

The following command enables forwarding of IP traffic for all VLANs with IP addresses:

```
enable ipforwarding
```

The following command enables forwarding of IP broadcast traffic for a VLAN named *accounting*:

```
enable ipforwarding broadcast vlan accounting
```

History

This command was available in ExtremeWare 2.0.

This command was modified in an ExtremeWare IP Technology Services Release based on ExtremeWare v6.1.8b12 to support MPLS modules.

Platform Availability

This command is available on all platforms.

enable ip-option loose-source-route

```
enable ip-option loose-source-route
```

Description

Enables the loose source route IP option.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

None.

Example

The following command enables the loose source route IP option:

```
enable ip-option loose-source-route
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on "*i*" series platforms.

enable ip-option record-route

```
enable ip-option record-route
```

Description

Enables the record route IP option.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

None.

Example

The following command enables the record route IP option:

```
enable ip-option record-route
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on "*i*" series platforms.

enable ip-option record-timestamp

```
enable ip-option record-timestamp
```

Description

Enables the record timestamp IP option.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

None.

Example

The following command enables the record timestamp IP option:

```
enable ip-option record-timestamp
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on "*i*" series platforms.

enable ip-option strict-source-route

```
enable ip-option strict-source-route
```

Description

Enables the strict source route IP option.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

None.

Example

The following command enables the strict source route IP option:

```
enable ip-option strict-source-route
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on "*i*" series platforms.

enable ip-option use-router-alert

```
enable ip-option use-router-alert
```

Description

Enables the generation of the router alert IP option.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

None.

Example

The following command enables generation of the router alert IP option:

```
enable ip-option use-router-alert
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on "*i*" series platforms.

enable iproute sharing

```
enable iproute sharing
```

Description

Enables load sharing if multiple routes to the same destination are available. When multiple routes to the same destination are available, load sharing can be enabled to distribute the traffic to multiple destination gateways. Only paths with the same lowest cost is will be shared.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

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.

Configure static routes and/or OSPF as you would normally. ExtremeWare supports unlimited route sharing across static routes and up to 12 ECMP routes for OSPF.

Route sharing is useful only in instances where you are constrained for bandwidth. This is typically not the case using Extreme switches. Using route sharing makes router troubleshooting more difficult because of the complexity in predicting the path over which the traffic will travel.

Example

The following command enables load sharing for multiple routes:

```
enable iproute sharing
```

History

This command was available in ExtremeWare 2.0.

This command was modified in ExtremeWare 6.2.2 to allow support of up to 12 ECMP routes for OSPF.

Platform Availability

This command is available on all platforms.

enable irdp

```
enable irdp {vlan <name>}
```

Description

Enables the generation of ICMP router advertisement messages on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

If no optional argument is specified, all the IP interfaces are affected.

Example

The following command enables IRDP on VLAN *accounting*:

```
enable irdp vlan accounting
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

enable loopback-mode vlan

```
enable loopback-mode vlan [<name> | all]
```

Description

Allows a VLAN to be placed in the UP state without an external active port. This allows (disallows) the VLANs routing interface to become active.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

N / A.

Usage Guidelines

Use this command to specify a stable interface as a source interface for routing protocols. This decreases the possibility of route flapping, which can disrupt connectivity.

Example

The following command allows the VLAN *accounting* to be placed in the UP state without an external active port:

```
enable loopback-mode vlan accounting
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

enable multinetting

```
enable multinetting
```

Description

Enables IP multinetting on the system.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The multinetting features requires the user to change the MAC FDB aging timer to be at least 3000 seconds on the switch. This command will automatically change the FDB timer to 3000 seconds if it is shorter than 3000 seconds.

Example

The following command enables multinetting on the system:

```
enable multinetting
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

enable subvlan-proxy-arp vlan

```
enable subvlan-proxy-arp vlan [<super-vlan name> | all]
```

Description

Enables the automatic entry of sub-VLAN information in the proxy ARP table.

Syntax Description

super-vlan name	Specifies a super-VLAN name.
all	Specifies all VLANs.

Default

Enabled.

Usage Guidelines

To facilitate communication between sub-VLANs, by default, an entry is made in the IP ARP table of the super-VLAN that performs a proxy ARP function. This allows clients on one sub-VLAN to communicate with clients on another sub-VLAN. In certain circumstances, intra-sub-VLAN communication may not be desired for isolation reasons.

NOTE

The isolation option works for normal, dynamic, ARP-based client communication.

Example

The following command enables the automatic entry of sub-VLAN information in the proxy ARP table opf the super-VLAN *vsuper*:

```
enable subvlan-proxy-arp vlan vsuper
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

rtlookup

```
rtlookup [<ip address> | <hostname>]
```

Description

Performs a look-up in the route table to determine the best route to reach an IP address or host.

Syntax Description

hostname	Specifies a hostname.
ip address	Specifies an IP address.

Default

N/A.

Usage Guidelines

The output of the `rtlookup` command has been enhanced to include information about MPLS LSPs associated with the routes. The flags field displayed by this command has been enhanced to indicate the presence of MPLS next hops. An uppercase `L` indicates the presence of a direct LSP next hop for the route. A lowercase `l` indicates the presence of an indirect LSP next hope for the route.

An optional `mpls` keyword has been added to the `rtlookup` command. When the `mpls` keyword is specified, the information displayed is modified; some of the information normally displayed is omitted, and the LSP endpoint and outgoing MPLS label are displayed instead. The LSP endpoint is the IP address/prefix of the FEC associated with the LSP. The LSP endpoint matches the destination for direct LSPs and is a 32-bit prefix address of a proxy router for indirect LSPs.

Example

The following command performs a look up in the route table to determine the best way to reach the specified hostname:

```
rtlookup berkeley.edu
```

History

This command was available in ExtremeWare 6.1.

This command was modified in an ExtremeWare IP Technology Services Release based on ExtremeWare v6.1.8b12 to support MPLS modules.

Platform Availability

This command is available on "*i*" series platforms.

run ipfdb-check

```
run ipfdb-check [index <bucket> <entry> | <ip-address> {<ip-address>}]
{extended} {detail}
```

Description

Checks IP FDB entries for consistency.

Syntax Description

bucket	Specifies the bucket portion of the FDB hash index.
entry	Specifies the entry portion of the FDB hash index.
ip-address	Specifies an IP address. FDB entries with this IP address will be checked.
ip-address	Specifies a second IP address, for checking bi-directional entries.
extended	Enables OTP index checking in the MAC entry and VPST of the egress port.
detail	Specifies that more detailed debug information should be logged.

Default

N / A.

Usage Guidelines

The IP FDB error checking function logs the error count into the system log. Use the `detail` option to log more detailed debug information.

Example

The following command will do consistency checking on IP FDB entries for IP address 10.20.30.55:

```
run ipfdb-check 10.20.30.55
```

History

This command was first available in ExtremeWare 6.1.9

Platform Availability

This command is available on all "*i*" series platforms.

The `extended` option is available on the Black Diamond 6800 chassis-based system only.

show iparp

```
show iparp {<ip address> | <mac_address> | vlan <name> | permanent}
```

Description

Displays the IP Address Resolution Protocol (ARP) table. You can filter the display by IP address, MAC address, VLAN, or permanent entries.

Syntax Description

ip address	Specifies an IP address.
mac_address	Specifies a MAC address.
name	Specifies a VLAN name.
permanent	Specifies permanent entries.

Default

Show all entries.

Usage Guidelines

Displays the IP ARP table, including:

- IP address
- MAC address
- Aging timer value
- VLAN name, VLAN ID and port number
- Flags

Example

The following command displays the IP ARP table:

```
show iparp 10.1.1.5/24
```

History

This command was available in ExtremeWare 2.0.

This command was modified in ExtremeWare 6.2.1 to provide the MAC address option.

Platform Availability

This command is available on all platforms.

show iparp proxy

```
show iparp proxy {<ip address> {<mask>}}
```

Description

Displays the proxy ARP table.

Syntax Description

ip address	Specifies an IP address.
mask	Specifies a subnet mask.

Default

N/A.

Usage Guidelines

If no argument is specified, then all proxy ARP entries are displayed.

Example

The following command displays the proxy ARP table:

```
show iparp proxy 10.1.1.5/24
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

show ipconfig

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

Description

Displays configuration information for one or more VLANs.

Syntax Description

name	Specifies a VLAN name.
detail	Specifies to display global IP configuration information in the detailed format.

Default

N/A.

Usage Guidelines

If no VLAN information is specified, then global IP configuration is displayed. Otherwise, specific VLAN(s) information will be displayed. Global IP configuration information includes:

- IP address/netmask/etc.
- IP forwarding information / IP multicast forwarding information
- BOOTP configuration
- VLAN name and VLANID
- ICMP configuration (global)
- IGMP configuration (global)
- IRDP configuration (global)

Example

The following command displays configuration information on a VLAN named *accounting*:

```
show ipconfig vlan accounting
```

History

This command was available in ExtremeWare 2.0.

This command was modified in an ExtremeWare IP Technology Services Release based on ExtremeWare v6.1.8b12 to support MPLS modules.

Platform Availability

This command is available on all platforms.

show ipfdb

```
show ipfdb {<ip address> <netmask> | vlan <name>}
```

Description

Displays the contents of the IP forwarding database (FDB) table. Used for technical support purposes. If no option is specified, all IP FDB entries are displayed.

Syntax Description

ip address	Specifies an IP address.
netmask	Specifies a subnet mask.
name	Specifies a VLAN name.

Default

Default is to show all IP FDB entries.

Usage Guidelines

Displays IP FDB table content including:

Dest IP Addr	IP address
TblIdx	IP FDB hash index and entry number
MacIdx	MAC FDB hash index and entry number
Flag	Flags
FlowInfo	
MAC Address	Next hop router MAC address
VLAN	Egress VLAN ID
Port	Egress port number

Example

The following command displays the contents of the IP FDB table on a VLAN named *accounting*:

```
show ipfdb vlan accounting
```

```
Dest IP Addr  TblIdx MacIdx Flag FlowInfo   MAC Address     VLAN Port
--------------- ------ ------ ---- -------- ----------------- ---- ----
10.205.4.201    00C3.0 9C32.0          0000 00:E0:2B:04:DA:00 4000 1
10.205.4.200    01C3.0 9C32.0          0000 00:E0:2B:04:DA:00 4000 1
10.205.4.203    02C3.0 9C32.0          0000 00:E0:2B:04:DA:00 4000 1
10.205.4.202    03C3.0 9C32.0          0000 00:E0:2B:04:DA:00 4000 1
10.205.4.205    04C3.0 9C32.0          0000 00:E0:2B:04:DA:00 4000 1
10.0.5.0        050F.0 9C32.0          0000 00:E0:2B:04:DA:00 4000 1
10.205.4.204    05C3.0 9C32.0          0000 00:E0:2B:04:DA:00 4000 1
10.205.4.207    06C3.0 9C32.0          0000 00:E0:2B:04:DA:00 4000 1
10.205.4.206    07C3.0 9C32.0          0000 00:E0:2B:04:DA:00 4000 1
```

```
10.205.0.202    07C7.0 4646.0          0000 00:10:E3:1D:00:1E 4000 1
10.205.4.193    08C3.0 9C32.0          0000 00:E0:2B:04:DA:00 4000 1
10.205.4.192    09C3.0 9C32.0          0000 00:E0:2B:04:DA:00 4000 1
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

show iproute

```
show iproute {priority | vlan <vlan> | permanent | <ip address> <netmask> |
route-map |origin [direct | static | blackhole | rip | bootp | icmp |
ospf-intra | ospf-inter | ospf-as-external | ospf-extern1 | ospf-extern2]}
{mpls} {sorted}
```

Description

Displays the contents of the IP routing table or the route origin priority.

Syntax Description

priority	Specifies a route priority.
vlan	Specifies a VLAN name.
permanent	Specifies permanent routing.
ip address	Specifies an IP address.
netmask	Specifies a subnet mask.
route-map	Specifies display of route maps for direct, static, blackhole, RIP, BOOTP, ICMP, OSPF-intra, OSPF-inter, OSPF as External, OSPF External 1, and OPPF External 2 routing.
origin	Specifies a display of the route map origin.
mpls	Specifies to display MPLS information.
sorted	Specifies to sort the information displayed.

Default

N / A.

Usage Guidelines

Route maps for IP routing can be configured based on the route origin. When routes are added to the IP routing table from various sources, the route map configured for the origin of the route is applied to the route. After matching on specified characteristics, the characteristics for the route can be modified using the route maps. The characteristics that can be matched and modified are dependent on the origin of the route. Route maps for IP routing can be dynamically changed. In the case of direct and static route origins, the changes are reflected immediately. In the case of routes that are sourced from other origin, the changes are reflected within 30 seconds.

The output of the `show iproute` command has been enhanced to include information about MPLS LSPs associated with the routes. The flags field displayed by this command has been enhanced to indicate the presence of MPLS next hops. An uppercase `L` indicates the presence of a direct LSP next hop for the route. A lowercase `l` indicates the presence of an indirect LSP next hope for the route.

An optional `mpls` keyword has been added to the `show iproute` command. When the `mpls` keyword is specified, the information displayed is modified; some of the information normally displayed is omitted, and the LSP endpoint and outgoing MPLS label are displayed instead. The LSP endpoint is the IP address / prefix of the FEC associated with the LSP. The LSP endpoint matches the destination for direct LSPs and is a 32-bit prefix address of a proxy router for indirect LSPs.

The mpls keyword only applies to some of the options available on the show iproute command. The mpls keyword is ignored when specified in conjunction with the following options:

- priority
- route-map
- summary

Example

The following command displays detailed information about all IP routing:

```
show iproute detail
```

Following is the output from this command:

```
Destination: 10.10.121.111/30
  Gateway: 10.10.121.201     VLAN   : helium          Origin : *d
  Metric : 1                   Flags  : U------u-       Time   : 13:15:26:49
  Use    : 14409               M-Use  : 0               Acct-1 : 0

 Destination: 10.11.166.112/29
  Gateway: 10.17.0.1          VLAN   : helium          Origin : *be
  Metric : 2                   Flags  : UG-----um       Time   : 01:11:23:49
  Use    : 0                   M-Use  : 0               Acct-1 : 0

 Destination: 10.13.105.112/29
  Gateway: 10.11.110.123     VLAN   : helium          Origin : *be
  Metric : 2                   Flags  : UG-----um       Time   : 00:29:09:23
  Use    : 0                   M-Use  : 0               Acct-1 :
```

History

This command was available in ExtremeWare 2.0.

This command was modified in an ExtremeWare IP Technology Services Release based on ExtremeWare v6.1.8b12 to support MPLS modules.

This command was modified to include a timestamp in ExtremeWare 6.2.2.

Platform Availability

This command is available on all platforms.

show ipstats

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

Description

Displays IP statistics for the CPU for the switch or for a particular VLAN.

Syntax Description

name	Specifies a VLAN name.

Default

N / A.

Usage Guidelines

This command only shows statistics of the CPU-handled packets. Not all packets are handled by the CPU.

The fields displayed in the `show ipstats` command are defined in Table 17 though Table 21.

Table 17: Global IP Statistics Field Definitions

Field	Definition
InReceives	Total number of incoming IP packets processed by the CPU.
InUnicast	Total number of unicast IP packets processed by the CPU.
InBcast	Total number of broadcast IP packets processed by the CPU.
InMcast	Total number of multicast IP packets processed by the CPU.
InHdrEr	Total number of packets with an IP Header Error forwarded to the CPU.
Bad vers	Total number of packets with a version other than IP v4 in the IP version field.
Bad chksum	Total number of packets with a bad IP checksum forwarded to the CPU.
Short pkt	IP packets that are too short.
Short hdr	IP packets with a header that is too short.
Bad hdrlen	IP packets with a header length that is less than the length specified.
Bad length	IP packets with a length less than that of the header.
InDelivers	IP packets passed to upper layer protocols.
Bad Proto	IP packets with unknown (not standard) upper layer protocol.
OutRequest	IP packets sent from upper layers to the IP stack.
OutDiscard	IP packets that are discarded due to lack of buffer space or the router interface being down, or broadcast packets with broadcast forwarding disabled.
OutNoRoute	IP packets with no route to the destination.
Forwards	ForwardOK and Fwd Err aggregate count.
ForwardOK	Total number of IP packets forwarded correctly.
Fwd Err	Total number of IP packets that cannot be forwarded.

Table 17: Global IP Statistics Field Definitions (continued)

Field	Definition
NoFwding	Aggregate number of IP packets not forwarded due to errors.
Redirects	IP packets forwarded on the same network.
No route	Not used.
Bad TTL	IP packets with a bad time-to-live.
Bad MC TTL	IP packets with a bad multicast time-to-live.
Bad IPdest	IP packets with an address that does not comply with the IP v4 standard.
Blackhole	IP packets with a destination that is a blackhole entry.
Output err	Not used. This is the same as Fwd Err.
MartianSrc	IP packets with an invalid source address.

Table 18: Global ICMP Statistics Field Definitions

Field	Definition
OutResp	Echo replies sent from the CPU.
OutError	Redirect from broadcast or multicast source addresses.
InBadcode	Incoming ICMP packets with an invalid CODE value.
InTooshort	Incoming ICMP packets that are too short.
Bad chksum	Incoming ICMP packets with checksum errors.
In Badlen	Incoming ICMP packets with length errors.
echo reply (In/Out):	ICMP “echo reply” packets that are received and transmitted.
destination unreachable (In/Out):	ICMP packets with destination unreachable that are received and transmitted.
port unreachable (In/Out):	ICMP packets with port unreachable that are received and transmitted.
echo (In/Out):	ICMP echo packets that are received and transmitted.

Table 19: Global IGMP Statistics Field Definitions

Field	Definition
Out Query	Number of IGMP query messages sent by the router.
Out Report	Number of reports sent on an active multicast route interface for reserved multicast addresses and for regular IGMP reports forwarded by the query router.
Out Leave	Number of IGMP out leave messages forwarded for IP multicast router interfaces.
In Query	Number of IGMP query messages received.
In Report	Number of IGMP report messages received (mostly from hosts).
In Leave	Number of IGMP leave messages received (mostly from hosts).
In Error	Number of IGMP packets with bad header fields or checksum failures.

Table 20: DHCP/BOOTP Statistics Field Definitions

Field	Definition
Received to server	Number of DHCP packets forwarded to server.
Received to client	Number of DHCP packets received by clients.
Requests relayed	Number of DHCP request packets relayed.
Responses relayed	Number of DHCP response packets relayed.
DHCP Discover	Number of DHCP Discover messages.
DHCP Offer	Number of DHCP Offer messages.
DHCP Request	Number of DHCP Request messages.
DHCP Decline	Number of DHCP Decline responses.
DHCP Ack	Number of DHCP Ack responses.
DHCP NAck	Number of DHCP NAck responses.
DHCP Release	Number of DHCP Release instances.
DHCP Inform	Not used.

Table 21: Router Interface Statistics Field Definitions

Field	Definition
Packets IN/OUT	Total number of IP packets received or transmitted on a VLAN router interface.
Octets IN/OUT	Total number of octets received or transmitted on a VLAN router interface.
Mcast packets IN/OUT	Total number of multicast packets received or transmitted on a VLAN router interface.
Bcast packets IN/OUT	Total number of broadcast packets received or transmitted on a VLAN router interface.
Errors IN/OUT	Total number of IP packets with errors received or transmitted on a VLAN router interface.
Discards IN/OUT	Total number of IP packets that cannot travel up to the CPU due to lack of buffer space.
Unknown Protocols IN/OUT	Total number of IP packets with unknown upper layer protocols received by the router interface.

Example

The following command displays IP statistics for the VLAN *accounting*:

```
show ipstats vlan accounting
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

show udp-profile

```
show udp-profile {<profile_name>}
```

Description

Displays the UDP profile information.

Syntax Description

profile_name	Specifies a UDP profile name.

Default

N/A.

Usage Guidelines

Displays the following information:

- Profile names
- Input rules of UDP port, destination IP address, or VLAN
- Source VLANs to which the profile is applied.

Example

The following command displays the UDP profile information for the UPD profile named *backbone*:

```
show udp-profile backbone
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

unconfig icmp

```
unconfig icmp
```

Description

Resets all ICMP settings to the default values.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command resets all ICMP settings to the default values.

```
unconfig icmp
```

History

This command was available in ExtremeWare 2.0.

This command was modified in an ExtremeWare IP Technology Services Release based on ExtremeWare v6.1.8b12 to support MPLS modules.

Platform Availability

This command is available on all platforms.

unconfig iparp

```
unconfig iparp
```

Description

Resets IP ARP timeout, IP ARP max-entries, and IP ARP max-pending-entries to their default values.

Syntax Description

This command has no arguments or variables.

Default

N/A

Usage Guidelines

None.

Example

The following command resets all IP ARP related settings to the default values:

```
unconfig iparp
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on "*i*" series platforms.

unconfig irdp

```
unconfig irdp
```

Description

Resets all router advertisement settings to the default values.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command resets all router advertisement settings to the default values.

```
unconfig irdp
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

unconfig udp-profile

```
unconfig udp-profile vlan [<name> | all]
```

Description

Removes the UDP-forwarding profile configuration for one or all VLANs.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all UDP profiles.

Default

N/A.

Usage Guidelines

None.

Example

The following command removes the UDP profile configuration from the VLAN *accounting*:

```
unconfig udp-profile vlan accounting
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

17 IGRP Commands

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 distributed among them. The cost of a route is described by a single metric.

OSPF allows parts of a networks 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 belonging to two or more areas. It is responsible for exchanging summary advertisements with other ABRs.
- Autonomous System Border Router (ASBR)—An ASBR acts as a gateway between OSPF and other routing protocols, or other autonomous systems.

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 LSDB remaining in use.

NOTE

Do not set the router ID to 0.0.0.0.

Routing Information Protocol (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.

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.

A new version of RIP, called RIP version 2 (RIPv2), expands the functionality of RIP version 1 to include:

- Variable-Length Subnet Masks (VLSMs)
- Next-hop addresses
- Support for next-hop addresses allows for optimization of routes in certain environments
- Multicasting

If you are using RIP with supernetting/Classless Inter-Domain Routing (CIDR), you must use RIPv2 only, and RIP route aggregation must be turned off.

config ospf cost

```
config ospf [area <areaid> | vlan [<name> | all]] cost [automatic |
<number>]]
```

Description

Configures the cost metric of one or all interface(s).

Syntax Description

areaid	Specifies an OSPF area.
name	Specifies a VLAN name.
all	Specifies all VLANs.
automatic	Specifies to determine the advertised cost from the OSPF metric table.
number	Specifies the cost metric.

Default

The default cost is 1.

Usage Guidelines

None.

Example

The following command configures the cost metric of the VLAN *accounting*:

```
config ospf area 0.0.0.6 vlan accounting cost 10
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config ospf priority

```
config ospf [area <areaid> | vlan [<name> | all]] priority <priority>
```

Description

Configures the priority used in the designated router-election algorithm for one or all OSPF interface(s) for all the interfaces within the area.

Syntax Description

areaid	Specifies an OSPF area.
name	Specifies a VLAN name.
all	Specifies all VLANs.
priority	Specifies a priority range. The range is 0 through 255.

Default

The default setting is 1.

Usage Guidelines

The range is 0 through 255, and the default setting is 1. Setting the value to 0 ensures that the router is never selected as the designated router or backup designated router.

Example

The following command sets the switch to not be selected as the designated router:

```
config ospf area 1.2.3.4 priority 0
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config ospf virtual-link authentication password

```
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>]
```

Description

Specifies the authentication password (up to eight characters) or Message Digest 5 (MD5) key for one or all interfaces in an area.

Syntax Description

name	Specifies a VLAN name.
areaid	Specifies an OSPF area.
routerid	Specifies a router interface number.
password	Specifies an authentication password (up to 8 ASCII characters).
md5-key_id	Specifies a Message Digest 5 key, from 0-255.
md5_key	Specifies a numeric value from 0-65,536. Can also be alphanumeric
none	Disables authentication.

Default

N / A.

Usage Guidelines

The md5_key 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.

Example

The following command configures MD5 authentication on the VLAN *subnet_26*:

```
config ospf vlan subnet_26 authentication md5 32 test
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config ospf timer

```
config ospf [vlan <name> | area <areaid> | virtual-link <routerid>
<areaid>] timer <retransmit interval> <transit delay> <hello interval>
<dead interval>
```

Description

Configures the timers for one interface or all interfaces in the same OSPF area.

Syntax Description

name	Specifies a VLAN name.
areaid	Specifies an OSPF area.
routerid	Specifies a router interface number.
retransmit interval	Specifies a retransmit interval in seconds. The range is 0 - 3,600 seconds.
transit delay	Specifies a transit delay in seconds. The range is 0 - 3,600 seconds.
hello interval	Specifies the hello interval in seconds. The range is 1 - 65,535 seconds.
dead interval	Specifies the dead interval in seconds. The range is 1 - 2,147,483,647 seconds.

Default

- retransmit interval — Default: 5
- transit delay — Default: 1
- hello interval — Default: 10
- dead interval — Default: 40

Usage Guidelines

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.

Example

The following command sets the timers on the virtual link in area 0.0.0.2:

```
config ospf virtual-link 6.6.6.6 0.0.0.2 timer 10 1 20 200
```

History

This command was available in ExtremeWare 2.0.

The syntax was modified in ExtremeWare 6.2.2.

Platform Availability

This command is available on all platforms.

config ospf add virtual-link

```
config ospf add virtual-link <routerid> <areaid>
```

Description

Adds a virtual link connected to another ABR.

Syntax Description

routerid	Specifies a router interface number.
areaid	Specifies an OSPF area.

Default

N / A.

Usage Guidelines

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. Specify the following:

- routerid — Far-end router interface number.
- areaid — Transit area used for connecting the two end-points. The transit area cannot have the IP address 0.0.0.0. the transit area cannot be a stub area or an NSSA.

Example

The following command configures a virtual link between the two interfaces:

```
config ospf add virtual-link 10.1.2.1 10.1.0.0
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config ospf add vlan area

```
config ospf add vlan [<name> | all] area <areaid> {passive}
```

Description

Enables OSPF on one or all VLANs (router interfaces).

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.
areaid	Specifies the area to which the VLAN is assigned.
passive	Specifies to stop sending and receiving hello packets on this interface.

Default

Disabled.

Usage Guidelines

When the peer LSR is also an Extreme switch, the following options are available for ensuring that an OSPF route is advertised for the tunnel endpoint IP address:

- A route is advertised when OSPF is enabled on the VLAN to which the IP address is assigned (using the `config ospf add vlan` command on the peer switch).
- A route is advertised when the peer switch is configured to distribute direct routes into the OSPF domain (via the `enable ospf export direct` command). The export option should be used when the tunnel LSP needs to cross OSPF area boundaries or when the Extreme Standby Routing Protocol (ESRP) is enabled on the VLAN to which the IP address is assigned.

In either case, LDP must be configured to advertise label mappings for direct routing interfaces.

In some configurations, you may want to enable loopback mode on the VLAN to which the tunnel endpoint IP address is assigned. One situation where loopback mode may be useful is when multiple physical interfaces, associated with different VLANs, are connected to the MPLS backbone. In this case, use of loopback-mode can provide redundancy by enabling TLS traffic to continue even when the physical interfaces associated with the tunnel endpoint IP address VLAN fail.

Example

The following command enables OSPF on a VLAN named *accounting*:

```
config ospf add vlan accounting area 0.0.0.1
```

History

This command was available in ExtremeWare 2.0. This command was modified in an ExtremeWare IP Technology Services Release based on ExtremeWare v6.1.8b12 to support MPLS modules.

Platform Availability

This command is available on all platforms.

config ospf add vlan area link-type

```
config ospf add vlan [<name> | all] area <areaid> link-type [auto |
broadcast | point-to-point] {passive}
```

Description

Configures the OSPF link type.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.
areaid	Specifies the area to which the VLAN is assigned.
auto	Specifies to automatically determine the OSPF link type based on the interface type.
broadcast	Specifies a broadcast link, such as Ethernet. Routers must elect a DR and a BDR during synchronization.
point-to-point	Specifies a point-to-point link type, such as PPP.
passive	Specifies to stop sending and receiving packets on this interface.

Default

Auto.

Usage Guidelines

The passive parameter indicates that the router only synchronizes and listens, and does not originate or send any new information on the interface.

Example

The following command configures the OSPF link type as automatic on a VLAN named *accounting*:

```
config ospf add vlan accounting area 0.0.0.1 link-type auto
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on "*i*" series platforms.

config ospf area external-filter

```
config ospf area <area_id> external-filter [<access_profile> |none]
```

Description

Configures an external filter policy.

Syntax Description

area_id	Specifies the OSPF target area.
access_profile	Specifies an access profile.
none	Specifies not to apply an external filter.

Default

N / A.

Usage Guidelines

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.

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.

Using the none mode specifies that no external filter is applied.

Example

The following command configures an external filter policy from the access profile *nosales*:

```
config ospf area 1.2.3.4 external-filter nosales
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

config ospf area interarea-filter

```
config ospf area <area_id> interarea-filter [<access_profile> | none]
```

Description

Configures a global inter-area filter policy.

Syntax Description

area_id	Specifies the OSPF target area.
access_profile	Specifies an access profile.
none	Specifies not to apply an interarea filter.

Default

N / A.

Usage Guidelines

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.

Example

The following command configures an inter-area filter policy from the access profile *nosales*:

```
config ospf area 0.0.0.6 interarea-filter nosales
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

config ospf area add range

```
config ospf area <areaid> add range <ipaddress> <mask> [advertise |
noadvertise] {type-3 | type-7}
```

Description

Configures a range of IP addresses in an OSPF area to be aggregated.

Syntax Description

areaid	Specifies an OSPF area.
ipaddress	Specifies an IP address
mask	Specifies a subnet mask.
advertise	Specifies to advertise the aggregated range of IP addresses.
noadvertise	Specifies not to advertise the aggregated range of IP addresses.
type-3	Specifies type 3 LSA, summary LSA.
type-7	Specifies type 7 LSA, NSSA external LSA.

Default

N/A.

Usage Guidelines

If advertised, the aggregated IP range is exported as a single LSA by the ABR.

Example

The following command is used to summarize a certain range of IP addresses within an area and export them out as a single address:

```
config ospf area 1.2.3.4 add range 10.1.2.0/24 advertise type-3
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config ospf area delete range

```
config ospf area <areaid> delete range <ipaddress> <mask>
```

Description

Deletes a range of aggregated IP addresses in an OSPF area.

Syntax Description

areaid	Specifies an OSPF area.
ipaddress	Specifies an IP address.
mask	Specifies a subnet mask.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes an aggregated IP address range:

```
config ospf area 1.2.3.4 delete range 10.1.2.0/24
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config ospf area normal

```
config ospf area <areaid> normal
```

Description

Configures an OSFP area as a normal area.

Syntax Description

areaid	Specifies an OSPF area.

Default

Normal.

Usage Guidelines

A normal area is an area that is not any of the following:

- Stub area
- NSSA

Virtual links can be configured through normal areas. External routes can be distributed into normal areas.

Example

The following command configures an OSPF area as a normal area:

```
config ospf area 10.1.0.0 normal
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config ospf area nssa stub-default-cost

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

Description

Configures an OSPF area as an NSSA.

Syntax Description

areaid	Specifies an OSPF area.
summary	Specifies that type-3 can be propagated into the area.
nosummary	Specifies that type-3 cannot be propagated into the area.
cost	Specifies a cost metric.
translate	Specifies whether type-7 LSAs are translated into type-5 LSAs.

Default

N / A.

Usage Guidelines

NSSAs are similar to the existing OSPF stub area configuration option, but have the following two additional capabilities:

- External originating from an ASBR connected to the NSSA can be advertised within the NSSA.
- External originating from the NSSA can be propagated to other areas, including the backbone area.

When configuring an OSPF area as an NSSA, the translate option 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.

Example

The following command configures an OSPF area as an NSSA:

```
config ospf area 10.1.1.0 nssa summary stub-default-cost 10 translate
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

config ospf area stub stub-default-cost

```
config ospf area <areaid> stub [summary | nosummary] stub-default-cost
<cost>
```

Description

Configures an OSPF area as a stub area.

Syntax Description

areaid	Specifies an OSPF area.
summary	Specifies that type-3 can be propagated into the area.
nosummary	Specifies that type-3 cannot be propagated into the area.
cost	Specifies a cost metric.

Default

N / A.

Usage Guidelines

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 and computation requirements on OSPF routers.

Example

The following command configures an OSPF area as a stub area:

```
config ospf area 0.0.0.6 stub nosummary stub-default-cost 10
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config ospf asbr-filter

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

Description

Configures a route filter for all ASBR routers.

Syntax Description

access_profile	Specifies an access profile.
none	Specifies not to apply an ASBR filter.

Default

N / A.

Usage Guidelines

For switches configured to support RIP, BGP, VIP 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.

Example

The following command configures a route filter for all routes OSPF exports from RIP or other sources:

```
config ospf asbr-filter subnet25-filter
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

config ospf ase-limit

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

Description

Configures the AS-external LSA limit and overflow duration associated with OSPF database overflow handling.

Syntax Description

number	Specifies the number of external routes that can be held on a link-state database.
seconds	Specifies a duration for which the system has to remain in the overflow state.

Default

The default for limit is 0, which indicates that there is no limit.

The default for timeout is 0, which indicates that once the router goes into overflow state, it stays there indefinitely.

Usage Guidelines

This command is only valid on an ASBR.

Example

The following command configures the AS-external LSA limit and overflow duration:

```
config ospf ase-limit 50000 timeout 1800
```

History

This command was available in ExtremeWare 6.1.5.

Platform Availability

This command is available on "*i*" series platforms.

config ospf ase-summary add

```
config ospf ase-summary add <ipaddress> <mask> cost <cost> {<tag_number>}
```

Description

Aggregates AS-external routes in a specified address range.

Syntax Description

ipaddress	Specifies an IP address.
mask	Specifies a subnet mask.
cost	Specifies a metric that will be given to the summarized route.
tag_number	Specifies an OSPF external route tag.

Default

N / A.

Usage Guidelines

This command is only valid on an ASBR.

Example

The following command summarizes AS-external routes:

```
config ospf ase-summary add 175.1.0.0/16 cost 10
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

config ospf ase-summary delete

```
config ospf ase-summary delete <ipaddress> <mask>
```

Description

Deletes an aggregated OSPF external route.

Syntax Description

ipaddress	Specifies an IP address.
mask	Specifies a subnet mask.

Default

N / A.

Usage Guidelines

This command is only valid on an ASBR.

Example

The following command deletes the aggregated AS-external route:

```
config ospf ase-summary delete 175.1.0.0/16
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on "*i*" series platforms.

config ospf delete virtual-link

```
config ospf delete virtual-link <routerid> <areaid>
```

Description

Removes a virtual link.

Syntax Description

routerid	Specifies a router interface number.
areaid	Specifies an OSPF area.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes a virtual link:

```
config ospf delete virtual-link 10.1.2.1 10.1.0.0
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config ospf delete vlan

```
config ospf delete vlan [<name> | all]
```

Description

Disables OSPF on one or all VLANs (router interfaces).

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables OSPF on VLAN *accounting*:

```
config ospf delete vlan accounting
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config ospf direct-filter

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

Description

Configures a route filter for direct routes.

Syntax Description

access_profile	Specifies an access profile.
none	Specifies not to apply a direct filter.

Default

N / A.

Usage Guidelines

If none is specified, all direct routes are exported if ospf export direct is enabled.

In versions of ExtremeWare before release 6.0, direct routes corresponding to the interfaces on which RIP was enabled were exported into OSPF as part of RIP routes, using the command enable ospf export rip. Using ExtremeWare 6.0 and above, you must configure ExtremeWare to export these direct routes to OSPF. You can use an access profile to filter unnecessary direct routes.

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.

Example

The following command configures a route filter for direct routes based on the access profile *nosales*:

```
config ospf direct-filter nosales
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

config ospf lsa-batching-timer

```
config ospf lsa-batching-timer <timer_value>
```

Description

Configures the OSPF LSA batching interval.

Syntax Description

timer_value	Specifies a time in seconds.

Default

The default setting is 30 seconds.

Usage Guidelines

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.

Example

The following command configures the OSPF LSA batching timer to a value of 100 seconds:

```
config ospf lsa-batching-timer 100
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

config ospf metric-table

```
config ospf metric-table 10m <cost> 100m <cost> 1g <cost>
```

Description

Configures the automatic interface costs for 10 Mbps, 100 Mbps, and 1 Gbps interfaces.

Syntax Description

cost	Specifies the interface cost for 10 Mbps interfaces.
cost	Specifies the interface cost for 100 Mbps interfaces.
cost	Specifies the interface cost for 1 Gbps interfaces.

Default

- 10 Mbps — The default cost is 10.
- 100 Mbps — The default cost is 5.
- 1 Gbps — The default cost is 4.

Usage Guidelines

None.

Example

The following command configures the automatic interface costs for 10 Mbps, 100 Mbps, and 1 Gbps interfaces:

```
config ospf metric-table 10m 20 100m 10 1g 2
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

config ospf routerid

```
config ospf routerid [automatic | <routerid>]
```

Description

Configures the OSPF router ID. If automatic is specified, the switch uses the largest IP interface address as the OSPF router ID.

Syntax Description

automatic	Specifies to use automatic addressing.
routerid	Specifies a router address.

Default

Automatic.

Usage Guidelines

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.

NOTE

Do not set the router ID to 0.0.0.0.

The implementation of the config ospf routerid command has been augmented to support automatic advertisement of a label mapping for the OSPF router ID. A label is advertised for the OSPF router ID regardless of whether OSPF distributes a route for the router ID IP address in its router LSA.

To support the use of indirect LSPs, Extreme LSRs automatically advertise a label mapping for a /32 LSP to its OSPF router ID (configured using the config ospf routerid command).

Example

The following command sets the router ID:

```
config ospf routerid 10.1.6.1
```

History

This command was available in ExtremeWare 2.0.

This command was modified in an ExtremeWare IP Technology Services Release based on ExtremeWare v6.1.8b12 to support MPLS modules.

Platform Availability

This command is available on all platforms.

config ospf spf-hold-time

```
config ospf spf-hold-time <seconds>
```

Description

Configures the minimum number of seconds between Shortest Path First (SPF) recalculations.

Syntax Description

seconds	Specifies a time in seconds.

Default

3 seconds.

Usage Guidelines

None.

Example

The following command configures the minimum number of seconds between Shortest Path First (SPF) recalculations:

```
config ospf spf-hold-time 6
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

config ospf vlan area

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

Description

Associates a VLAN (router interface) with an OSPF area. By default, all router interfaces are associated with area 0.0.0.0.

Syntax Description

name	Specifies a VLAN name.
areaid	Specifies an OSPF area.

Default

Area 0.0.0.0

Usage Guidelines

Any OSPF network that contains more than one area is required to have an area configured as area 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, 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, by default you must assign it to an area.

Example

The following command associates the VLAN *accounting* with an OSPF area:

```
config ospf vlan accounting area 0.0.0.6
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config ospf vlan neighbor add

```
config ospf vlan <name> neighbor add <ipaddress>
```

Description

Configures the IP address of a point-to-point neighbor.

Syntax Description

name	Specifies a VLAN name.
ipaddress	Specifies an IP address.

Default

N/A.

Usage Guidelines

None.

Example

The following command configures the IP address of a point-to-point neighbor:

```
config ospf vlan accounting neighbor add 10.0.0.1
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on "*i*" series platforms.

config ospf vlan neighbor delete

```
config ospf vlan <name> neighbor delete <ipaddress>
```

Description

Deletes the IP address of a point-to-point neighbor.

Syntax Description

name	Specifies a VLAN name.
ipaddress	Specifies an IP address.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes the IP address of a point-to-point neighbor:

```
config ospf vlan accounting neighbor delete 10.0.0.1
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on "*i*" series platforms.

config ospf vlan timer

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

Description

Configures the OSPF wait interval.

Syntax Description

vlan	Specifies a VLAN name.
rxmtinterval	Specifies the length of time that the router waits before retransmitting an LSA that is not acknowledged.
transitdelay	Specifies the length of time it takes to transmit an LSA packet over the interface.
hellointerval	Specifies the interval at which routers send hello packets.
routerdeadinterval	Specifies the interval after which a neighboring router is declared down due to the fact that hello packets are no longer received from the neighbor.
waitinterval	Specifies the interval between the interface coming up and the election of the DR and BDR.

Default

- rxmtinterval--5 seconds.
- hellointerval--10 seconds.
- routerdeadinterval--40 seconds.

Usage Guidelines

Specify the following:

- rxmtinterval—If you set an interval that is too short, unnecessary retransmissions will result.
- transitdelay—The transit delay must be greater than 0.
- hellointerval—Smaller times allow routers to discover each other more quickly, but also increase network traffic.
- routerdeadinterval—This interval should be a multiple of the hello interval.
- waitinterval—This interval is required by the OSPF standard to be equal to the routerdeadinterval. Under some circumstances, setting the waitinterval 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 hellointerval. The default value is equal to the routerdeadinterval.

Example

The following command configures the OSPF wait interval on the VLAN *accounting*:

```
config ospf vlan accounting timer 10 15 20 60 60
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on "*i*" series platforms.

config rip add vlan

```
config rip add vlan [<name> | all]
```

Description

Configures RIP on an IP interface.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

All. If no VLAN is specified, then all is assumed.

Usage Guidelines

When an IP interface is created, RIP configuration is disabled on the interface by default. When the RIP interface is disabled, the parameters are not reset to default automatically.

Example

The following command configures RIP on the VLAN *finance*:

```
config rip add finance
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config rip delete vlan

```
config rip delete vlan [<name> | all]
```

Description

Disables RIP on an IP interface.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

All. If no VLAN is specified, then all is assumed.

Usage Guidelines

When an IP interface is created, RIP configuration is disabled on the interface by default. When the RIP interface is disabled, the parameters are not reset to default automatically.

Example

The following command deletes RIP on a VLAN named *finance*:

```
config rip delete finance
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config rip garbagetime

```
config rip garbagetime {<seconds>}
```

Description

Configures the RIP garbage time.

Syntax Description

seconds	Specifies a time in seconds.

Default

120 seconds.

Usage Guidelines

None.

Example

The following command configures the RIP garbage time to have a 60-second delay:

```
config rip garbagetime 60
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config rip routetimeout

```
config rip routetimeout {<seconds>}
```

Description

Configures the route timeout period.

Syntax Description

seconds	Specifies a time in seconds.

Default

180 seconds.

Usage Guidelines

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.

Example

The following example sets the route timeout period to 120 seconds:

```
config rip routetimeout 120
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config rip rxmode

```
config rip rxmode [none | v1only | v2only | any] {vlan <name>}
```

Description

Changes the RIP receive mode for one or more VLANs.

Syntax Description

none	Specifies to drop all received RIP packets.
v1only	Specifies to accept only RIP version 1 format packets.
v2only	Specifies to accept only RIP version 2 format packets.
any	Specifies to accept RIP version 1 and RIP version 2 packets.
name	Specifies to apply settings to specific VLAN name.

Default

Any.

Usage Guidelines

If no VLAN is specified, the setting is applied to all VLANs.

Example

The following command configures the receive mode for the VLAN *finance* to accept only RIP version 1 format packets:

```
config rip rxmode v1only finance
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config rip txmode

```
config rip txmode [none | v1only | v1comp | v2only] {vlan <name>}
```

Description

Changes the RIP transmission mode for one or more VLANs.

Syntax Description

none	Specifies to not transmit any packets on this interface.
v1only	Specifies to transmit RIP version 1 format packets to the broadcast address.
v1comp	Specifies to transmit RIP version 2 format packets to the broadcast address.
v2only	Specifies to transmit RIP version 2 format packets to the RIP multicast address.
name	Specifies to apply settings to a specific VLAN name.

Default

v2only.

Usage Guidelines

If no VLAN is specified, the setting is applied to all VLANs.

Example

The following command configures the transmit mode for the VLAN *finance* to transmit version 2 format packets to the broadcast address:

```
config rip txmode v1comp finance
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config rip updatetime

```
config rip updatetime {<seconds>}
```

Description

Specifies the time interval in seconds within which RIP sends update packets.

Syntax Description

seconds	Specifies a time in seconds.

Default

30 seconds.

Usage Guidelines

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*). The timer granularity is 10 seconds.

Example

The following command sets the update timer to 60 seconds:

```
config rip updatetime 60
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config rip vlan cost

```
config rip vlan [<name> | all] cost <number>
```

Description

Configures the cost (metric) of the interface.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.
number	Specifies a cost metric.

Default

The default setting is 1.

Usage Guidelines

None.

Example

The following command configures the cost for the VLAN *finance* to a metric of 3:

```
config rip vlan finance cost 3
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

config rip vlan export-filter

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

Description

Configures RIP to suppress certain routes when performing route advertisements.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.
access_profile	Specifies an access-profile name.
none	Specifies to check the access profile for permit and deny attributes.

Default

N / A.

Usage Guidelines

Use an access profile to determine trusted RIP router neighbors for the VLAN on the switch running RIP.

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. When 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.

Example

The following command uses the access profile *nosales* to determine which RIP routes are advertised into the VLAN *backbone*:

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

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

config rip vlan import-filter

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

Description

Configures RIP to ignore certain routes received from its neighbor.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.
access_profile	Specifies an access-profile name.
none	Specifies to check the access profile for permit and deny attributes.

Default

N / A.

Usage Guidelines

Configures an import filter policy, which uses 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.

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. When 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.

Example

The following command configures the VLAN *backbone* to accept selected routes from the access profile *nosales*:

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

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

config rip vlan trusted-gateway

```
config rip vlan [<name> | all] trusted-gateway [<access_profile> | none]
```

Description

Configures a trusted neighbor policy, which uses an access profile to determine trusted RIP router neighbors for the VLAN on the switch running RIP.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.
access_profile	Specifies an access-profile name.
none	Specifies to check the access profile for permit and deny attributes.

Default

N/A.

Usage Guidelines

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. When 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.

Example

The following command configures RIP to use the access profile *nointernet* to determine from which RIP neighbor to receive (or reject) the routes to the VLAN *backbone*:

```
config rip vlan backbone trusted-gateway nointernet
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

create ospf area

```
create ospf area <areaid>
```

Description

Creates an OSPF area.

Syntax Description

areaid	Specifies an OSPF area.

Default

Area 0.0.0.0

Usage Guidelines

Area 0.0.0.0 does not need to be created. It exists by default.

Example

The following command creates an OSPF area:

```
create ospf area 1.2.3.4
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

delete ospf area

```
delete ospf area [<areaid> | all]
```

Description

Deletes an OSPF area.

Syntax Description

areaid	Specifies an OSPF area.
all	Specifies all areas.

Default

N/A.

Usage Guidelines

An OSPF area cannot be deleted if it has an associated interface.

Example

The following command deletes an OSPF area:

```
delete ospf area 1.2.3.4
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

disable ospf

```
disable ospf
```

Description

Disables the OSPF process for the router.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables the OSPF process for the router:

```
disable ospf
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

disable ospf capability opaque-lsa

```
disable ospf capability opaque-lsa
```

Description

Disables opaque LSAs across the entire system.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

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.

If your network uses opaque LSAs, all routers on your OSPF network should support opaque LSAs. Routers that do not support opaque LSAs do not store or flood them. At minimum a well-interconnected subsection of your OSPF network needs to support opaque LSAs to maintain reliability of their transmission.

On an OSPF broadcast network, the designated router (DR) must support opaque LSAs or none of the other routers on that broadcast network will reliably receive them. You can use the OSPF priority feature to give preference to an opaque-capable router, so that it becomes the elected DR.

For transmission to continue reliably across the network, the backup designated router (BDR) must also support opaque LSAs.

Example

The following command disables opaque LSAs across the entire system:

```
disable ospf capability opaque-lsa
```

History

This command was available in ExtremeWare 6.2.

Platform Availability

This command is available on "*i*" series platforms.

disable ospf export

```
disable ospf export [bgp | i-bgp | e-bgp] [cost <number> [ase-type-1 |
ase-type-2] {tag <number>} | <route_map>]
```

Description

Disables redistribution of routes to OSPF.

Syntax Description

bgp	Specifies BGP routing.
i-bgp	Specifies I-BGP routing.
e-bgp	Specifies E-BGP routing.
number	Specifies a cost metric.
ase-type-1	Specifies AS-external type 1 routes.
ase-type-2	Specifies AS-external type 2 routes.
number	Specifies a tag value.
route_map	Specifies a route map.

Default

The default tag number is 0. The default setting is disabled.

Usage Guidelines

After it is enabled, the OSPF router is considered to be an ASBR. Interface routes that correspond to the interface that has OSPF enabled are ignored.

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.

The cost metric is inserted for all BGP, VIP, 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. When a route map is associated with the export command, the route map is applied on every exported route. The exported routes can also be filtered using route maps.

Example

The following command disables OSPF to export BGP-related routes using LSAs to other OSPF routers:

```
disable ospf export bgp cost 1 ase-type-1 tag 0
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

disable ospf export direct

```
disable ospf export direct [cost <metric> [ase-type-1 | ase-type-2] {tag
<number>} | <route_map>]
```

Description

Disables the redistribution of local interface (direct) routes into the OSPF domain. This will not export the loopback address of 127.0.0.1.

Syntax Description

metric	Specifies a cost metric.
ase-type-1	Specifies AS-external type 1 routes.
ase-type-2	Specifies AS-external type 2 routes.
number	Specifies a tag value.
route_map	Specifies a route map.

Default

The default tag number is 0. The default setting is disabled.

Usage Guidelines

After it is enabled, the OSPF router is considered to be an ASBR. Interface routes that correspond to the interface that has OSPF enabled are ignored.

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.

The cost metric is inserted for all BGP, VIP, 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. When a route map is associated with the export command, the route map is applied on every exported route. The exported routes can also be filtered using route maps.

Example

The following command disables the distribution of local interface (direct) routes into the OSPF domain:

```
disable ospf export direct cost 1 ase-type-1 tag 0
```

History

This command was first available in ExtremeWare 6.1.

This command was modified in an ExtremeWare IP Technology Services Release based on ExtremeWare v6.1.8b12 to support MPLS modules.

Platform Availability

This command is available on "*i*" series platforms.

disable ospf export rip

```
disable ospf export rip [cost <metric> [ase-type-1 | ase-type-2] {tag
<number>} | <route map>]
```

Description

Disables the redistribution of RIP to OSPF.

Syntax Description

metric	Specifies a cost metric.
ase-type-1	Specifies AS-external type 1 routes.
ase-type-2	Specifies AS-external type 2 routes.
number	Specifies a tag value.
route map	Specifies a route map name.

Default

The default tag number is 0. The default setting is disabled.

Usage Guidelines

After it is enabled, the OSPF router is considered to be an ASBR.

This command enables the exporting of RIP by way of LSA to other OSPF routers as AS-external type 1 or type 2 routes.

The cost metric is inserted for all BGP, VIP, RIP-learned, static, and direct routes injected into OSPF. 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.

When re-distributing RIP routes, you should turn off RIP aggregation unless you are expertly familiar with the possible consequences and impact. By default, new configurations of RIP using ExtremeWare 4.0 and above disable RIP aggregation. In previous ExtremeWare versions, RIP aggregation is enabled by default. This configuration is preserved when upgrading to ExtremeWare 4.0. Verify the configuration using the command `show rip`.

Example

The following command disables the exporting of RIP to OSPF:

```
disable ospf export rip cost 1 ase-type-1 tag 0
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

disable ospf export static

```
disable ospf export static [cost <metric> [ase-type-1 | ase-type-2] {tag
<number>} | <route_map>]
```

Description

Disables the redistribution of static routes to OSPF.

Syntax Description

metric	Specifies a cost metric.
ase-type-1	Specifies AS-external type 1 routes.
ase-type-2	Specifies AS-external type 2 routes.
number	Specifies a tag value.
route_map	Specifies a route map name.

Default

The default tag number is 0. The default setting is disabled.

Usage Guidelines

After it is enabled, the OSPF router is considered to be an ASBR.

This command enables the redistribution of static routes by way of LSA to other OSPF routers as AS-external type 1 or type 2 routes.

The cost metric is inserted for all BGP, VIP, RIP-learned, static, and direct routes injected into OSPF. 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.

Example

The following command disables the exporting of static routes to OSPF:

```
disable ospf export static cost 0 ase-type-1 tag 0
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

disable ospf export vip

```
disable ospf export vip [cost <metric> [ase-type-1 | ase-type-2] {tag
<number>} | <route_map>]
```

Description

Disables the redistribution of virtual IP addresses into the OSPF domain.

Syntax Description

metric	Specifies a cost metric.
ase-type-1	Specifies AS-external type 1 routes.
ase-type-2	Specifies AS-external type 2 routes.
number	Specifies a tag value.
route_map	Specifies a route map.

Default

The default tag number is 0. The default setting is disabled.

Usage Guidelines

After it is enabled, the OSPF router is considered to be an ASBR.

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.

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 BGP, VIP, 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 sued for selective insertion. When a route map is associated with the export command, the route map is applied on every exported route. The exported routes can also be filtered using route maps.

Example

The following command disables the redistribution of virtual IP addresses into the OSPF domain:

```
disable ospf export vip cost 0 ase-type-1 tag 0
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

disable rip

```
disable rip
```

Description

Disables RIP for the whole router.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

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

Example

The following command disables RIP for the whole router:

```
disable rip
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

disable rip aggregation

```
disable rip aggregation
```

Description

Disables the RIP aggregation of subnet information on a RIP version 2 (RIPv2) interface.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

The enable (disable) rip aggregation command enables (disables) the RIP aggregation of subnet information on an interface configured to send RIPv1 or RIPv2-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 RIPv1.
- If aggregation is disabled, subnet routes are never aggregated, even when crossing a class boundary.

Example

The following command disables RIP aggregation on the interface:

```
disable rip aggregation
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

disable rip export metric

```
disable rip export [static | direct | ospf | ospf-intra | opst-inter |
ospf-extern1 | ospf-extern2 | vip] metric <metric> {tag <number>}
```

Description

Disables RIP to redistribute routes from other routing functions.

Syntax Description

static	Specifies static routes.
direct	Specifies interface routes (only interfaces that have IP forwarding enabled are exported).
ospf	Specifies all OSPF routes.
ospf-intra	Specifies OSPF-intra area routes.
ospf-inter	Specifies OSPF-inter area routes.
ospf-extern1	Specifies OSPF external route type 1.
ospf-extern2	Specifies OSPF external route type 2.
vip	Specifies VIP routes.
metric	Specifies the `metric` range, from 0-15. If set to 0, RIP uses the route metric obtained from the route origin.
number	Specifies a tag number.

Default

Disabled.

Usage Guidelines

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 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.

Example

The following command disables RIP to redistribute routes from all OSPF routes:

```
disable rip export ospf
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

disable rip exportstatic

```
disable rip exportstatic
```

Description

Enables the redistribution of static routes.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Static routes are manually entered into the routing table. Static routes are used to reach networks not advertised by routers. You can configure 64 static unicast routes. 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. 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.

Example

The following command disables the redistribution of static routes:

```
disable rip exportstatic
```

History

This command was removed in ExtremeWare 6.0.

Platform Availability

This command is available on "*i*" series platforms.

disable rip originate-default cost

```
disable rip originate-default {always} cost <metric> {tag<number>}
```

Description

Unconfigures a default route to be advertised by RIP if no other default route is advertised. If always is specified, RIP always advertises the default route to its neighbors. If always is not specified, RIP adds a default route if a reachable default route is not in the route table.

Syntax Description

always	Specifies to always advertise the default route.
metric	Specifies a cost metric.
number	Specifies a tag number.

Default

N/A

Usage Guidelines

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.

Example

The following command unconfigures a default route to be advertised by RIP if no other default route is advertised:

```
disable rip originate-default cost 0
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

disable rip poisonreverse

```
disable rip poisonreverse
```

Description

Disables poison reverse algorithm for RIP.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

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.

Example

The following command disables the split horizon with poison reverse algorithm for RIP:

```
disable rip poisonreverse
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

disable rip splithorizon

```
disable rip splithorizon
```

Description

Disables the split horizon algorithm for RIP.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

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.

Example

The following command disables the split horizon algorithm for RIP:

```
disable rip splithorizon
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

disable rip triggerupdate

```
disable rip triggerupdate
```

Description

Disables the trigger update mechanism. Triggered updates are a mechanism for immediately notifying a router's neighbors when the router adds or deletes routes or changes their metric.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

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.

Example

The following command disables the trigger update mechanism:

```
disable rip triggerupdate
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

enable ospf

```
enable ospf
```

Description

Enables the OSPF process for the router.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command enables the OSPF process for the router:

```
enable ospf
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

enable ospf capability opaque-lsa

```
enable ospf capability opaque-lsa
```

Description

Enables opaque LSAs across the entire system.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

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.

If your network uses opaque LSAs, all routers on your OSPF network should support opaque LSAs. Routers that do not support opaque LSAs do not store or flood them. At minimum a well-interconnected subsection of your OSPF network needs to support opaque LSAs to maintain reliability of their transmission.

On an OSPF broadcast network, the designated router (DR) must support opaque LSAs or none of the other routers on that broadcast network will reliably receive them. You can use the OSPF priority feature to give preference to an opaque-capable router, so that it becomes the elected DR.

For transmission to continue reliably across the network, the backup designated router (BDR) must also support opaque LSAs.

Example

The following command enables opaque LSAs across the entire system:

```
enable ospf capability opaque-lsa
```

History

This command was available in ExtremeWare 6.2.

Platform Availability

This command is available on "*i*" series platforms.

enable ospf export

```
enable ospf export [bgp | i-bgp | e-bgp] [cost <number> [ase-type-1 |
ase-type-2] {tag <number>} | <route_map>]
```

Description

Enables redistribution of routes to OSPF.

Syntax Description

bgp	Specifies BGP routing.
i-bgp	Specifies I-BGP routing.
e-bgp	Specifies E-BGP routing.
number	Specifies a cost metric.
ase-type-1	Specifies AS-external type 1 routes.
ase-type-2	Specifies AS-external type 2 routes.
number	Specifies a tag value.
route_map	Specifies a route map.

Default

The default tag number is 0. The default setting is disabled.

Usage Guidelines

After it is enabled, the OSPF router is considered to be an ASBR. Interface routes that correspond to the interface that has OSPF enabled are ignored.

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.

The cost metric is inserted for all BGP, VIP, 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. When a route map is associated with the export command, the route map is applied on every exported route. The exported routes can also be filtered using route maps.

Example

The following command enables OSPF to export BGP-related routes using LSAs to other OSPF routers:

```
enable ospf export bgp cost 1 ase-type-1 tag 0
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

enable ospf export direct

```
enable ospf export direct [cost <metric> [ase-type-1 | ase-type-2] {tag
<number>} | <route_map>]
```

Description

Enables the redistribution of local interface (direct) routes into the OSPF domain. This will not export the loopback address of 127.0.0.1.

Syntax Description

metric	Specifies a cost metric.
ase-type-1	Specifies AS-external type 1 routes.
ase-type-2	Specifies AS-external type 2 routes.
number	Specifies a tag value.
route_map	Specifies a route map.

Default

The default tag number is 0. The default setting is disabled.

Usage Guidelines

After it is enabled, the OSPF router is considered to be an ASBR. Interface routes that correspond to the interface that has OSPF enabled are ignored.

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.

The cost metric is inserted for all BGP, VIP, 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. When a route map is associated with the export command, the route map is applied on every exported route. The exported routes can also be filtered using route maps.

Example

The following command enables the distribution of local interface (direct) routes into the OSPF domain:

```
enable ospf export direct cost 1 ase-type-1 tag 0
```

History

This command was first available in ExtremeWare 6.1.

This command was modified in an ExtremeWare IP Technology Services Release based on ExtremeWare v6.1.8b12 to support MPLS modules.

Platform Availability

This command is available on "*i*" series platforms.

enable ospf export rip

```
enable ospf export rip [cost <metric> [ase-type-1 | ase-type-2] {tag
<number>} | <route map>]
```

Description

Enables the redistribution of RIP to OSPF.

Syntax Description

metric	Specifies a cost metric.
ase-type-1	Specifies AS-external type 1 routes.
ase-type-2	Specifies AS-external type 2 routes.
number	Specifies a tag value.
route map	Specifies a route map name.

Default

The default tag number is 0. The default setting is disabled.

Usage Guidelines

After it is enabled, the OSPF router is considered to be an ASBR.

This command enables the exporting of RIP by way of LSA to other OSPF routers as AS-external type 1 or type 2 routes.

The cost metric is inserted for all BGP, VIP, RIP-learned, static, and direct routes injected into OSPF. 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.

When re-distributing RIP routes, you should turn off RIP aggregation unless you are expertly familiar with the possible consequences and impact. By default, new configurations of RIP using ExtremeWare 4.0 and above disable RIP aggregation. In previous ExtremeWare versions, RIP aggregation is enabled by default. This configuration is preserved when upgrading to ExtremeWare 4.0. Verify the configuration using the command `show rip`.

Example

The following command enables the exporting of RIP to OSPF:

```
enable ospf export rip cost 1 ase-type-1 tag 0
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

enable ospf export static

```
enable ospf export static [cost <metric> [ase-type-1 | ase-type-2] {tag
<number>} | <route_map>]
```

Description

Enables the redistribution of static routes to OSPF.

Syntax Description

metric	Specifies a cost metric.
ase-type-1	Specifies AS-external type 1 routes.
ase-type-2	Specifies AS-external type 2 routes.
number	Specifies a tag value.
route_map	Specifies a route map name.

Default

The default tag number is 0. The default setting is disabled.

Usage Guidelines

After it is enabled, the OSPF router is considered to be an ASBR.

This command enables the redistribution of static routes by way of LSA to other OSPF routers as AS-external type 1 or type 2 routes.

The cost metric is inserted for all BGP, VIP, RIP-learned, static, and direct routes injected into OSPF. 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.

Example

The following command enables the exporting of static routes to OSPF:

```
enable ospf export static cost 0 ase-type-1 tag 0
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

enable ospf export vip

```
enable ospf export vip [cost <metric> [ase-type-1 | ase-type-2] {tag
<number>} | <route_map>]
```

Description

Enables the redistribution of virtual IP addresses into the OSPF domain.

Syntax Description

metric	Specifies a cost metric.
ase-type-1	Specifies AS-external type 1 routes.
ase-type-2	Specifies AS-external type 2 routes.
number	Specifies a tag value.
route_map	Specifies a route map.

Default

The default tag number is 0. The default setting is disabled.

Usage Guidelines

After it is enabled, the OSPF router is considered to be an ASBR.

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.

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 BGP, VIP, 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 sued for selective insertion. When a route map is associated with the export command, the route map is applied on every exported route. The exported routes can also be filtered using route maps.

Example

The following command enables the redistribution of virtual IP addresses into the OSPF domain:

```
enable ospf export vip cost 0 ase-type-1 tag 0
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

enable ospf originate-default cost

```
enable ospf originate-default {always} cost <metric> [ase-type-1 |
ase-type-2] {tag <number>}
```

Description

Enables 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.

Syntax Description

always	Specifies for OSPF to always advertise the default route.
metric	Specifies a cost metric.
ase-type-1	Specifies AS-external type 1 routes.
ase-type-2	Specifies AS-external type 2 routes.
number	Specifies a tag value.

Default

N/A.

Usage Guidelines

If always is specified, OSPF always advertises the default route. If always is not specified, OSPF adds the default LSA if a reachable default route is in the route table.

Example

The following command generates a default external type-1 LSA:

```
enable ospf originate-default cost 1 ase-type-1 tag 0
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

enable rip

```
enable rip
```

Description

Enables RIP for the whole router.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

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

Example

The following command enables RIP for the whole router:

```
enable rip
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

enable rip aggregation

```
enable rip aggregation
```

Description

Enables the RIP aggregation of subnet information on a RIP version 2 (RIPv2) interface.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

The enable (disable) rip aggregation command enables (disables) the RIP aggregation of subnet information on an interface configured to send RIPv1 or RIPv2-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 RIPv1.
- If aggregation is disabled, subnet routes are never aggregated, even when crossing a class boundary.

Example

The following command enables RIP aggregation on the interface:

```
enable rip aggregation
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

enable rip export metric

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

Description

Enables RIP to redistribute routes from other routing functions.

Syntax Description

static	Specifies static routes.
direct	Specifies interface routes (only interfaces that have IP forwarding enabled are exported).
ospf	Specifies all OSPF routes.
ospf-intra	Specifies OSPF-intra area routes.
ospf-inter	Specifies OSPF-inter area routes.
ospf-extern1	Specifies OSPF external route type 1.
ospf-extern2	Specifies OSPF external route type 2.
vip	Specifies VIP routes.
metric	Specifies the `metric` range, from 0-15. If set to 0, RIP uses the route metric obtained from the route origin.
number	Specifies a tag number.

Default

Disabled.

Usage Guidelines

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 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.

Example

The following command enables RIP to redistribute routes from all OSPF routes:

```
enable rip export ospf
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

enable rip exportstatic

```
enable rip exportstatic
```

Description

Enables the redistribution of static routes.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

Static routes are manually entered into the routing table. Static routes are used to reach networks not advertised by routers. You can configure 64 static unicast routes. 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. 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.

Example

The following command enables the redistribution of static routes:

```
enable rip exportstatic
```

History

This command was removed in ExtremeWare 6.0.

Platform Availability

This command is available on "*i*" series platforms.

enable rip originate-default cost

```
enable rip originate-default {always} cost <metric> {tag<number>}
```

Description

Configures a default route to be advertised by RIP if no other default route is advertised. If always is specified, RIP always advertises the default route to its neighbors. If always is not specified, RIP adds a default route if a reachable default route is not in the route table.

Syntax Description

always	Specifies to always advertise the default route.
metric	Specifies a cost metric.
number	Specifies a tag number.

Default

N/A.

Usage Guidelines

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.

Example

The following command configures a default route to be advertised by RIP if no other default route is advertised:

```
enable rip originate-default cost 0
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

enable rip poisonreverse

```
enable rip poisonreverse
```

Description

Enables poison reverse algorithm for RIP.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

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.

Example

The following command enables the split horizon with poison reverse algorithm for RIP:

```
enable rip poisonreverse
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

enable rip splithorizon

```
enable rip splithorizon
```

Description

Enables the split horizon algorithm for RIP.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

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.

Example

The following command enables the split horizon algorithm for RIP:

```
enable rip splithorizon
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

enable rip triggerupdate

```
enable rip triggerupdate
```

Description

Enables the trigger update mechanism. Triggered updates are a mechanism for immediately notifying a router's neighbors when the router adds or deletes routes or changes their metric.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

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.

Example

The following command enables the trigger update mechanism:

```
enable rip triggerupdate
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

show ospf

```
show ospf
```

Description

Displays global OSPF information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays global OSPF information:

```
show ospf
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

show ospf area

```
show ospf area <areaid>
```

Description

Displays information about a particular OSPF area.

Syntax Description

areaid	Specifies an OSPF area.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays information about OSPF area 1.2.3.4:

```
show ospf area 1.2.3.4
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

show ospf area detail

```
show ospf area detail
```

Description

Displays information about all OSPF areas.

Syntax Description

detail	Specifies to display the information in detailed format.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays information about all OSPF areas:

```
show ospf area detail
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on "*i*" series platforms.

show ospf ase-summary

```
show ospf ase-summary
```

Description

Displays the OSPF external route aggregation configuration.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the OSPF external route aggregation configuration:

```
show ospf ase-summary
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on "*i*" series platforms.

show ospf interfaces detail

```
show ospf interfaces detail
```

Description

Displays detailed information about all OSPF interfaces.

Syntax Description

detail	Specifies to display the information in detailed format.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays information about all OSPF interfaces:

```
show ospf interfaces
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on "*i*" series platforms.

show ospf interfaces

```
show ospf interfaces {vlan <name> | area <areaid>}
```

Description

Displays information about one or all OSPF interfaces.

Syntax Description

name	Specifies a VLAN name.
areaid	Specifies an OSPF area.

Default

If no argument is specified, all OSPF interfaces are displayed.

Usage Guidelines

None.

Example

The following command displays information about one or all OSPF interfaces on the VLAN *accounting*:

```
show ospf interfaces vlan accounting
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

show ospf lsdb area lstype

```
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>]]
```

Description

Displays a table of the current LSDB.

Syntax Description

detail	Specifies to display all fields of matching LSAs in a multi-line format.
summary	Specifies to display several important fields of matching LSAs, one line per LSA.
stats	Specifies to display the number of matching LSAs, but not any of their contents.
all	Specifies all OSPF areas.
areaid	Specifies an OSPF area.
all	Specifies to display all LSA types.
as-external	Specifies to display type-5 LSAs.
external-type7	Specifies to display type-7 LSAs.
network	Specifies to display type-2 LSAs.
router	Specifies a remote router ID.
summary-asb	Specifies to display type-4 LSAs.
summary-net	Specifies to display type-3 LSAs.
id	Specifies an LS ID.
id	Specifies a router ID.

Default

Display in summary format.

Usage Guidelines

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.

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.

You can filter the display using either the area ID, the remote router ID, or the link-state ID. The default setting is `all` with no detail. If detail is specified, each entry includes complete LSA information.

Example

The following command displays all areas and all types in a summary format:

```
show ospf lsdb
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

show ospf virtual-link

```
show ospf virtual-link {<areaid> <routerid>}
```

Description

Displays virtual link information about a particular router or all routers.

Syntax Description

areaid	Specifies an OSPF area.
routerid	Specifies a router interface number.

Default

N / A.

Usage Guidelines

areaid — Transit area used for connecting the two end-points. The transit area cannot have the IP address 0.0.0.0.

Example

The following command displays virtual link information about a particular router:

```
show ospf virtual link 1.2.3.4 10.1.6.1
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

show rip

```
show rip {detail}
```

Description

Displays RIP specific configuration and statistics for all VLANs.

Syntax Description

detail	Specifies to display the information in detailed format.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays RIP specific configuration and statistics for all VLANs:

```
show rip
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

show rip stat

```
show rip stat {detail}
```

Description

Displays RIP-specific statistics for all VLANs.

Syntax Description

detail	Specifies to display the information in detailed format.

Default

All.

Usage Guidelines

Statistics include the following per interface:

- Packets transmitted
- Packets received
- Bad packets received
- Bad routes received
- Number of RIP peers
- Peer information

Example

The following command displays RIP-specific statistics for all VLANs:

```
show rip stat
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

show rip stat vlan

```
show rip stat vlan <name>
```

Description

Displays RIP specific statistics for a VLAN.

Syntax Description

name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays RIP specific statistics for the VLAN *accounting*:

```
show rip stat accounting
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

show rip vlan

```
show rip vlan <name>
```

Description

Displays RIP configuration and statistics for a VLAN.

Syntax Description

name	Specifies a VLAN name.

Default

All.

Usage Guidelines

None.

Example

The following command displays RIP configuration and statistics for the VLAN *accounting*:

```
show rip vlan accounting
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

unconfig ospf

```
unconfig ospf {vlan <name> | area <areaid>}
```

Description

Resets one or all OSPF interfaces to the default settings.

Syntax Description

name	Specifies a VLAN name.
areaid	Specifies an OSPF area.

Default

N/A.

Usage Guidelines

None.

Example

The following command resets the OSPF interfaces to the default settings on the VLAN *accounting*:

```
unconfig ospf accounting
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on "*i*" series platforms.

unconfig rip

```
unconfig rip {vlan <name>}
```

Description

Resets all RIP parameters to the default VLAN.

Syntax Description

name	Specifies a VLAN name.

Default

All.

Usage Guidelines

Does not change the enable/disable state of the RIP settings.

Example

The following command deletes RIP configuration from the VLAN *finance*:

```
unconfig rip finance
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

18 PoS Commands

The Packet over SONET (PoS) modules are I/O modules for the BlackDiamond 6800 series chassis-based system. These modules connect a BlackDiamond 6800 series switch to the SONET infrastructure used by metropolitan area service providers and operators of server co-location networks. (The BlackDiamond 6800 series switch is a chassis-based switch designed to be placed in the core of your network.)

Two key applications for the PoS modules are: interconnecting metropolitan area networks across the SONET network infrastructure, and interconnecting server co-location network sites directly using SONET links.

In the first application, the metropolitan area network service provider can build service network sites in various cities, then use PoS modules in a BlackDiamond 6800 series switch to connect those cities to a carrier's SONET infrastructure.

In the second application, operators of server co-location networks can use PoS modules in BlackDiamond 6800 series switches to create a SONET-based connection between server co-location sites. The result is that their network is simpler to manage, and problems can be isolated and resolved more expediently.

This chapter documents the PoS command set. Some commands are new for the PoS modules; other commands have been enhanced to support the PoS modules.

NOTE

Support for Packet over SONET modules is included in an ExtremeWare IP Services Technology Release, currently based on ExtremeWare v6.1.8b12. Later versions of ExtremeWare (6.1.9 or 6.2) currently do not support PoS modules.

config aps

```
config aps <group#> [nonrevert | revert <minutes>]
```

Description

Configures APS operation in either nonrevertive or revertive switching mode.

Syntax Description

group#	Specifies the APS group# to which the command applies.
nonrevert	Specifies nonrevertive switching mode when traffic is active on the protection line and the working line becomes operational.
revert	Specifies revertive switching mode when traffic is active on the protection line and the working line becomes operational.
minutes	Specifies the wait-to-restore (WTR) period in minutes.

Default

The default mode is nonrevertive switching.

Usage Guidelines

You specify the group# parameter, which is an integer in the range [1-65535]. The group# identifies the APS group that the configuration command applies to. The default mode is nonrevertive switching. This parameter determines what action should be taken when traffic is active on the protection line and the working line becomes operational. In revertive mode, traffic will automatically be switched from the protection line to the working line, after the user-defined wait-to-restore (WTR) period, which may be specified via the minutes parameter. The WTR period is intended to prevent frequent switches due to intermittent errors on the working line; service is restored only if no errors are detected on the working line during the WTR period. The minutes parameter is an integer in the range [0-12]. Conversely, in nonrevertive mode, traffic will remain on the protection line (until either manual intervention or a failure on the protection line forces a switch back to the working line). This parameter is only applicable to SONET ports performing the protection line function.

Example

The following command configures an APS operation on group 1001 in revertive switching mode for 5 minutes:

```
config APS 1001 revert 5
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config aps add

```
config aps <group#> add <port> [working | protection <ip address>]
```

Description

Adds a SONET port to an APS group.

Syntax Description

group#	Specifies the APS group# to which the command applies.
port	Specifies the SONET port number to be added to the APS group.
working	Specifies that the port is the working line.
protection	Specifies that the port is the protection line.
ip address	Specifies the IP address of the BlackDiamond switch where working line resides.

Default

By default, no ports are added to an APS group. Ports must be explicitly added using this command for proper APS operation.

Usage Guidelines

You specify the group# parameter, which is an integer in the range [1-65535]. The group# identifies the APS group that the port is to be added to. You also specify the port parameter, which identifies the SONET port that is to be added to the APS group. Additionally, you specify whether the port is designated as the working or protection line. Only one working line and one protection line can be added to an APS group. If the port is designated as the protection line, then you must also specify an IP address (ip address) of the BlackDiamond switch where the working line resides. This IP address is used to send APS control messages to the BlackDiamond switch containing the working line. It is recommended that the configured ip address be associated with an Ethernet VLAN that has loopback mode enabled (to minimize the impact of network outages on APS functionality). It is important that the network connecting working and protection switches always has sufficient bandwidth to support APS control transfers.

In routing configurations, the working line and the protection line should represent the same IP address from a neighboring PPP router's perspective. When the working line and protection line reside in the same BlackDiamond switch, this implies that both ports should be members of the same VLAN. The case where both the working line and the protection line for an APS group reside in the same BlackDiamond switch is the only situation where PPP's IP control protocol (IPCP) can be enabled on multiple SONET ports that are members of the same VLAN. In general, if IPCP is enabled on a SONET, then the port can only be a member of one VLAN, and no others ports can be members of that VLAN.

Example

The following command example adds port 1 of the module installed in slot 8 of the BlackDiamond switch to APS group 1001 as the working line:

```
config aps 1001 add 8:1 working
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config aps authenticate

```
config aps <group#> authenticate [off | on <string>]
```

Description

Configures authentication of APS control messages.

Syntax Description

group#	Specifies the APS group# to which the command applies.
off	Specifies that authentication is turned off.
on	Specifies authentication is turned on.
string	Specifies the authentication string used to validate the APS control frames received over an Ethernet link.

Default

The default setting is off.

Usage Guidelines

You specify the group# parameter, which is an integer in the range [1-65535]. The group# identifies the APS group that the authentication command applies to. You also specify whether authentication is to be turned off or turned on. If authentication is being enabled, a text authentication string must also be specified. This string can contain up to eight alphanumeric characters. If the working line and the protection line for an APS group reside in different BlackDiamond switches, then the same string must be configured at both BlackDiamond switches for authentication to work properly. The authentication string is used to validate APS control frames received over an Ethernet link. If authentication fails, the associated APS control frame is discarded.

Example

The following command example enables APS authentication for group 1001, with seer5dog as the authentication string:

```
config aps 1001 authenticate on seer5dog
```

History

This command was first available in ExtremeWare 6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config aps delete

```
config aps <group#> delete <port>
```

Description

Deletes a SONET port from an APS group.

Syntax Description

group#	Specifies the APS group# to which the command applies.
port	Specifies the port number.

Default

N / A.

Usage Guidelines

You specify the group# parameter, which is an integer in the range [1-65535]. The group# identifies the APS group that the port is to be deleted from. You also specify the port parameter, which identifies the SONET port that is to be deleted from the APS group. If you delete the working line from a group, it causes a switch to the protection line; however, if you delete an active protection line from a group, it does not initiate a switch to the working line.

Example

The following command example deletes port 1 of the module installed in slot 8 of the BlackDiamond switch from APS group 1001:

```
config aps 1001 delete 8:1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config aps force

```
config aps <group#> force [off | working | protection]
```

Description

Requests that an APS group be forced to use a specified line as the active line.

Syntax Description

group#	Specifies the APS group# to which the command applies.
off	Specifies that force is disabled.
working	Specifies that the APS group uses the working line as the active line.
protection	Specifies that the APS group uses the protection line as the active line.

Default

The default is `force off`.

Usage Guidelines

You specify the group# parameter, which is an integer in the range [1-65535]. The group# identifies the APS group that the force command applies to. When `force working` is specified, the command requests that the APS group uses the working line as the active line. Conversely, when `force protection` is specified, the command requests that the APS group uses the protection line as the active line. A forced switch is a high priority request. Only three events can override a forced switch request: (1) a `force off` command, (2) a `lockout on` command (that was either in effect before the force command or issued after the force command), or (3) a signal-fail condition on the protection line. This command is only applicable to SONET ports performing the protection line function. Additionally, the effects of this command are not preserved across a switch reboot.

Example

The following command example forces APS group `1001` to use the protection line as the active line:

```
config aps 1001 force protection
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config aps lockout

```
config aps <group#> lockout [off | on]
```

Description

Controls whether a switch to the protection line is locked out.

Syntax Description

group#	Specifies the APS group# to which the command applies.
off	Specifies that switches from the working line to the protection line are allowed.
on	Specifies that switches from the working line to the protection line are prohibited.

Default

The default is off.

Usage Guidelines

The group# identifies the APS group that the lockout command applies to. When lockout on is specified, switches from the working line to the protection line are prohibited, until you subsequently issue a **lockout off** command. The default is lockout off. This command is only applicable to SONET ports performing the protection line function. Additionally, the effects of this command are not preserved across a switch reboot.

Example

The following command example turns on lockout mode for APS group 1001:

```
config aps 1001 lockout on
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config aps manual

```
config aps <group#> manual [off | working | protection]
```

Description

Manually determines whether an APS group uses its working line or its protection line as the active line.

Syntax Description

group#	Specifies the APS group# to which the command applies.
off	Specifies that manual switching is disabled, and can be overridden.
working	Specifies that the APS group uses the working line as the active line.
protection	Specifies that the APS group uses the protection line as the active line.

Default

The default is `manual off`.

Usage Guidelines

You specify the group# parameter, which is an integer in the range [1-65535]. The group# identifies the APS group that the command applies to. When `manual working` is specified, the command requests that the APS group uses the working line as the active line. Conversely, when `manual protection` is specified, the command requests that the APS group uses the protection line as the active line. One potential use of the `manual working` command is to switch back to the working line after an error condition has cleared without waiting for the full wait-to-restore period to elapse. A manual switch is a lower priority request than a forced switch. events that can override a manual switch include: (1) a `manual off` command, (2) a `force working` or a force protection command, (3) a `lockout on` command, or (4) a signal-fail or signal degrade line condition. This command is only applicable to SONET ports performing the protection line function. Additionally, the effects of this command are not preserved across a switch reboot.

Example

The following command example configures APS group `1001` to use its working line as the active line:

```
config aps 1001 manual working
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config aps timers

```
config aps <group#> timers <seconds> <consecutive_misses>
```

Description

Sets the values of the timers used in the APS hello protocol that is exchanged between the working and protection switches for an APS group.

Syntax Description

group#	Specifies the APS group# to which the command applies.
seconds	Specifies the amount of time in seconds the protection switch waits between transmissions of hello packets to the working switch.
consecutive_misses	Specifies the time interval the protection switch will wait before assuming the working switch has failed.

Default

The default values are seconds = 1 and consecutive_misses = 5. These parameters are only applicable to SONET ports performing the protection line function.

Usage Guidelines

You specify the group# parameter, which is an integer in the range [1-65535]. The group# identifies the APS group that the configuration command applies to. The seconds parameter is an integer in the range [1-300] that specifies the amount of time the protection switch waits between transmissions of hello packets to the working switch. The consecutive_misses parameter is an integer in the range [1-100] that controls the time interval the protection switch will wait before assuming that the working switch has failed. If the working switch does not respond within consecutive_misses hello intervals, or (consecutive_misses * seconds) seconds, then the protection switch assumes that the working switch has failed and initiates a line switch.

Example

The following command example configures the timers for APS group 1001 to 1 second and 3 consecutive misses:

```
config aps 1001 timers 1 3
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config diffserv dscp-mapping ports

```
config diffserv dscp-mapping <input_codepoint>/<output_codepoint> ports
<portlist> {egress {no-congestion | congestion} | ingress}
```

Description

Configures a mapping between an input DiffServ code point (DSCP) and an associated output DSCP for the specified PoS ports.

Syntax Description

input_codepoint	Specifies one of the 64 possible DiffServ code point values as the input code point.
output_codepoint	Specifies one of the 64 possible DiffServ code point values as the output code point.
portlist	Specifies the port number(s).
egress	Applies the DSCP mapping to the egress direction.
no-congestion	Applies the DSCP mapping to the egress mapping table for the non-congested state.
congestion	Applies the DSCP mapping to the egress mapping table for the congested state.
ingress	Applies the DSCP mapping to the ingress direction.

Default

By default, all the tables are initialized such that DSCPs are not altered by the mapping operations; for example, an input DSCP value of n is always mapped to an output DSCP value of n. Additionally, `dscp-mapping` is performed without regard to whether `diffserv examination` is enabled on the port.

Usage Guidelines

Three DSCP mapping tables are supported per SONET port. One of the tables is used in the ingress direction and two are used for egress flows (onto the SONET link). The two egress tables are for the congested and non-congested states, as determined by the RED algorithm (in other words, the congested state is when the average queue length is greater than the minimum RED threshold). If RED is not enabled on the SONET port, then the egress congested-state mapping table is not used.

The tables are very simple. In the ingress direction, the input DSCP of a packet received from the SONET link is replaced with an output DSCP before the packet is forwarded. The replacement is straightforward; the input DSCP is used as an index into a 64-entry table that contains the output DSCPs associated with each of the input DSCP values. The operation is similar in the egress direction, with the DSCP mapping occurring before the packet is transmitted onto the SONET link(s). The mapping operation is performed after the packet has been assigned to a QoS profile. One potential use of the DSCP mapping capability is reconciliation of varying DiffServ policies at the boundary between autonomous systems (for example, at the boundary between two ISPs). The availability of different tables for the congested / non-congested states is useful for marking operations that increase the drop probability of packets during times of congestion, as discussed in the DiffServ assured forwarding (AF) RFC.

This command is currently only applicable to SONET ports. If the `no-congestion/congestion` keywords are omitted, the mapping is applied to the egress tables for both states. If the `egress/ingress` keywords are omitted, the mapping is assumed to apply to the egress direction, and a symmetrical mapping (with the `input_codepoint` and `output_codepoint` reversed) is automatically configured in the `ingress` direction.

Example

The following command example configures the congested-state mappings for DSCPs 10 (AF11):

```
config diffserv dscp-mapping 10/12 egress congestion
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config dot1q tagmapping ports

```
config dot1q tagmapping <input_vlanid/output_vlanid> ports <portlist>
{egress {priority <priority>} | ingress {priority <priority>}}
```

Description

Configures the VLAN tag mapping attributes for a PoS port.

Syntax Description

input_vlanid	Specifies VLAN ID of the input to be mapped.
output_vlanid	Specifies the VLAN ID of the output to be mapped.
portlist	Specifies the port number(s).
ingress	Indicates that the mapping is to be applied to input frames received from the PPP link.
egress	Indicates that the mapping is to be applied to input frames going to the PPP link.
priority	Allows you to set the 802.1p priority value.

Default

The default is to initialize the tables so the VLAN IDs are not altered by the mapping operations (for example, an input VLAN ID of n is always mapped to an output VLAN ID of n), and the frame priority is preserved.

Usage Guidelines

This command is only applicable when BCP is enabled on the port. Currently, the command is only supported for PoS ports. Two mapping tables are supported per PoS port. One of the tables is used in the egress direction and the other table is used in the ingress direction. Each of the tables enable an input VLAN ID to be mapped to an output VLAN ID, which can be useful in reconciling policy differences at customer/service provider boundaries. The egress keyword indicates that the mapping is to be applied to frames received from the switch backplane before transmission onto the PoS link(s). Conversely, the ingress keyword indicates that the mapping is to be applied to input frames received from the PoS link. The mappings are applied following classification to a QoS profile.

Frames containing the specified input_vlanid are altered such that the VLAN ID is set to the specified output_vlanid before the frame is forwarded. The tables also allow the option of preserving the 802.1p priority or overwriting the priority field with a configured value. The priority keyword indicates that the 802.1p priority field is to be set to the value of the priority parameter. Omission of the priority keyword indicates that the 802.1p priority of the frame is to be preserved. If the egress/ingress keywords are omitted, the specified mapping is applied to the egress direction, and a symmetrical mapping (with the input_vlanid and output_vlanid reversed) is automatically configured in the ingress direction. The input_vlanid and output_vlanid parameters are integers in the range [1-4095]. The priority parameter is an integer in the range [0-7].

Example

The following command configures the tagmapping attributes for input VLAN ID 30 and output VLAN ID 130 for port 1 of the module installed in slot 8 for the input frames from the PPP link:

```
config dot1q tagmapping 30/130 port 8:1 ingress
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config dot1q tagnesting ports

```
config dot1q tagnesting {<vlanid> | <vlanid_range>} [off | pop | push
<new_vlanid> {priority <priority>}] ports <portlist> {egress | ingress}
```

Description

Configures the VLAN tag nesting attributes for a PoS port. Currently, the command is only supported for PoS ports.

Syntax Description

vlanid	Specifies that the tag nesting will be performed on the frames containing the VLAN ID given.
vlanid_range	Specifies that the tag nesting will be performed on the frames containing VLAN ID values in the given range.
off	Disables tag nesting.
pop	Deletes a tag from the frame.
push	Adds a tag to the frame.
new_vlanid	Specifies the VLAN ID of the tag to be added or deleted from the frame.
priority	Allows you to set the 802.1p priority value.
portlist	Specifies the port number(s).
egress	Specifies that the tag operations are to be performed to the PPP link.
ingress	Specifies that the tag operations are to be performed from the PPP link.

Default

By default, tag nesting is off for all VLAN IDs. If the `egress/ingress` keywords are omitted, the direction defaults to `egress`. Additionally, if the `egress/ingress` keywords are omitted and a tag push operation is configured, a corresponding tag pop operation is automatically configured for the `ingress` direction. Similarly, if the `egress/ingress` keywords are omitted and tag nesting is configured off, it is disabled in both directions.

Usage Guidelines

The command provides support for nested 802.1Q tags by allowing a tag push/pop attribute to be associated with a VLAN ID. The push attribute indicates that a new tag is to be added to the frame, while the pop attribute indicates that the top-level tag is to be removed from the frame. The `push` keyword indicates that a new tag is to be added to frames containing the specified `vlanid` or one of the VLAN IDs in the specified `vlanid_range`. The syntax of the `vlanid_range` parameter is `start_vlanid-end_vlanid`. Omission of the `vlanid/vlanid_range` parameter indicates that the command settings should be applied to all VLAN IDs. For push operations, the new tag added to frames contains the specified `new_vlanid`.

The `pop` keyword indicates that the top-level tag is to be removed from frames when the tag contains any of the specified VLAN IDs. Tag operations may be performed in either `egress` (to the PoS link) or `ingress` directions.

When a new tag is pushed, an option is available to allow the 802.1p priority of the frame to be either preserved or set to a configured value. The priority keyword indicates that the 802.1p priority field is to

be set to the value of the priority parameter. Omission of the `priority` keyword indicates that the 802.1p priority of the frame is to be preserved. The `vlanid` parameters are integers in the range [1-4095]. The `priority` parameter is an integer in the range [0-7].

This command is only applicable when BCP is enabled on the port. Furthermore, tag push operations are applicable to egress frames only when the port is configured to transmit tagged frames for the associated VLAN. The tag-nesting operations are only applicable to `ingress` frames that contain a VLAN tag. The tag-nesting operations are applied after classification to a QoS profile. The default PPP MRU is sufficient for a single level of tag nesting (where the frame contains two VLAN tags) between two Extreme switches; jumbo frame support must be enabled if higher levels of VLAN tag nesting are needed.

The DiffServ/RED functions are not performed by PoS ports when frames contain nested tags (in other words, more than one tag).

Example

The following command adds VLAN `140` to the frame for port 1 of the module installed in slot 8 for input frames from the PPP link:

```
config dot1q tagnesting push 140 port 8:1 ingress
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config flowstats export add

```
config flowstats export {<group#>} add [<ip address> | <hostname>]
<udp_port>
```

Description

Configures the flow-collector devices to which NetFlow datagrams are exported.

Syntax Description

export <group#>	Specifies a particular export distribution group on a per-filter basis. Identifies the set of flow collector devices to which records for flows matching the filter are to be exported.
ip address	Specifies the IP address of the flow-collector destination.
hostname	Specifies the host name of the flow-collector destination.
udp_port	Specifies the UDP port number of the flow-collector destination.

Default

By default, no flow-collector destinations are configured.

Usage Guidelines

A flow-collector destination is identified by either an IP address and UDP port #, or by a hostname and UDP port #, to which NetFlow export datagrams are transmitted. The command allows flow-collector destinations to be added. Up to 8 flow-collector destinations can be configured for each group, and up to 32 groups can be defined per switch. The optional group# parameter, which identifies the specific group the destination is being configured for, is an integer in the range [1..32]. The group# defaults to 1 if the parameter is omitted. At least one flow-collector destination must be configured for NetFlow datagrams to be exported to a group.

When multiple flow-collectors are configured as members of the same group, the exported NetFlow datagrams are distributed across the available destinations. This NetFlow-distribution feature enables a scalable collection architecture that is able to accommodate high volumes of exported data. The distribution algorithm ensures that all the records for a given flow are exported to the same collector. The algorithm also ensures that flow records for both the ingress and egress directions of a TCP or UDP connection are exported to the same collector (when both flows traverse the same SONET link and both filters are configured to export to the same group).

Example

The following command adds a flow-collector destination of 10.1.1.88 for group 5 using UDP port 2025 to which NetFlow datagrams are exported:

```
config flowstat export 5 add 10.1.1.88 2025
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config flowstats export delete

```
config flowstats export {<group#>} delete [<ip address> | <hostname>]
<udp_port>
```

Description

Configures the flow-collector devices to which NetFlow datagrams are exported.

Syntax Description

export <group#>	Specifies a particular export distribution group on a per-filter basis. Identifies the set of flow collector devices to which records for flows matching the filter are to be exported.
ip address	Specifies the IP address of the flow-collector destination.
hostname	Specifies the host name of the flow-collector destination.
udp_port	Specifies the UDP port number of the flow-collector destination.

Default

N / A.

Usage Guidelines

A flow-collector destination is identified by either an IP address and UDP port #, or by a hostname and UDP port #, to which NetFlow export datagrams are transmitted. The command allows flow-collector destinations to be `deleted`. Up to 8 flow-collector destinations can be configured for each group, and up to 32 groups can be defined per switch. The optional `group#` parameter, which identifies the specific group the destination is being configured for, is an integer in the range [1..32]. The `group#` defaults to 1 if the parameter is omitted. At least one flow-collector destination must be configured for NetFlow datagrams to be exported to a group.

When multiple flow-collectors are configured as members of the same group, the exported NetFlow datagrams are distributed across the available destinations. This NetFlow-distribution feature enables a scalable collection architecture that is able to accommodate high volumes of exported data. The distribution algorithm ensures that all the records for a given flow are exported to the same collector. The algorithm also ensures that flow records for both the ingress and egress directions of a TCP or UDP connection are exported to the same collector (when both flows traverse the same SONET link and both filters are configured to export to the same group).

Example

The following command deletes a flow-collector destination of `10.1.1.88` for group `5` to which NetFlow datagrams are exported:

```
config flowstat export 5 delete 10.1.1.88 2025
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config flowstats filter ports

```
config flowstats filter <filter#> {aggregation} {export <group#>} ports
<portlist> [ingress | egress] <filterspec>
```

Description

Configures a flow record filter for the specified SONET ports.

Syntax Description

filter#	The `filter#` parameter is an integer in the range from 1 to 8 that operates with either the `ingress` or `egress` keyword to identify the filter that is being defined.
aggregation	Reduces the volume of exported data, use this optional keyword to maintain a single set of statistics for all the flows that match the specified filter.
export <group# tn	Specifies a particular export distribution group on a per-filter basis. Identifies the set of flow collector devices to which records for flows matching the filter are to be exported.
portlist	Specifies the port number(s).
ingress	Use this keyword to specify that the filter being defined in the command is one of the eight filters to be applied to ingress flows.
egress	Use this keyword to specify that the filter being defined in the command is one of the eight filters to be applied to egress flows.

filterspec	Each filter is defined using a *value/filtermask* pair for each of the five components in the following sequence:
	{des tination IP address, source IP address,
	destination port number, source port number, protocol}
	in the form:
	`[{dest-ip <ipaddress_value/ipaddress_filtermask>}`
	`{source-ip <ipaddress_value/ipaddress_filtermask>}`
	`{dest-port <port_value/port_filtermask>}`
	`{source-port <port_value/port_filtermask>}`
	`{protocol <protocol_value/protocol_filtermask>} |`
	`match-all-flows | match-no-flows]`
	The `ipaddress_filtermask`, `port_filtermask`, and `protocol_filtermask` parameters are configured using hexadecimal notation.
	You can also use either the `match-all-flows` keyword or the `match-no-flows` keyword in place of settings for the five components. The `match-all-flows` keyword adjusts the *value/filtermask* settings for all the components to 0/0 such that the filter matches any flow. The `match-no-flows` keyword adjusts the *value/filtermask* settings for all the components such that the filter does not match any flow.

Default

By default, filter#1 is configured to `match-all-flows`, and the remaining filters are configured to `match-no-flows`. The `group#` defaults to 1 if the parameter is omitted.

Usage Guidelines

The command allows a port to be configured to selectively maintain statistics for only those flows that match the specified filters. Sixteen filters are supported for each port, eight filters for ingress flows and another eight filters for egress flows. The `filter#` parameter and either the `ingress` or `egress` keyword are specified to identify the filter that is being configured. The `filter#` parameter is an integer in the range [1..8]. The filters are comprised of a value/filtermask pair for each component of the {destination IP address, source IP address, destination port number, source port number, protocol} 5-tuple. Conceptually, the filters work by ANDing the contents of each 5-tuple component of a forwarded flow with the associated masks from filter#1. Statistics are maintained if the results of the AND operations match the configured filter values for all fields of the 5-tuple. If there is no match, then the operation is repeated for filter#2, and so on. If there is no match for any of the filters, then statistics are not maintained for the flow. Filters for any/all of the 5-tuple components can be configured with a single command.

The `filterspec` parameter also supports the `match-all-flows` and `match-no-flows` keywords. The `match-all-flows` keyword adjusts the settings such that the filter matches any flow (that is, the value/filtermask pairs are set to 0/0 for all the 5-tuple components), while the `match-no-flows` keyword adjusts the settings such that the filter does not match any flow.

The optional `aggregation` keyword may be used to indicate that a single set of statistics is to be maintained for all the flows that match the filter, which can substantially reduce the volume of exported data. A particular export distribution group may also be specified on a filter-basis. The `group#` parameter identifies the set of collector devices that records for flows matching the filter are to be exported to.

Example

The following command example configures a filter to collect statistics on ingress flows destined for `192.168.1.1` from the `192.169.0.0/16` subnet with a destination port of `80` using protocol `6`:

```
config flowstats filter 1 export 1 ports all ingress
   dest-ip 192.168.1.1/FFFFFFFF source-ip 192.169.0.0/FFFF0000
   dest-port 80/FFFF source-port 0/0 protocol 6/FF
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config flowstats source ipaddress

```
config flowstats source ipaddress <ip address>
```

Description

Configures the IP address that is to be used as the source IP address for NetFlow datagrams to be exported.

Syntax Description

ip address	Specifies the source IP address to be used as the source for NetFlow datagrams to be exported.

Default

Normal.

Usage Guidelines

No NetFlow datagrams will be exported until the source `ip address` is configured. Flow-collector devices may use the source IP address of received NetFlow datagrams to identify the switch that sent the information. It is recommended that the configured `ip address` be associated with a VLAN that has loopback mode enabled.

Example

The following command example specifies that the IP address `192.168.100.1` is to be used as the source IP address for exported NetFlow datagrams:

```
config flowstats source ipaddress 192.168.100.1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20 for the PoS module only.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config ports tunnel hdlc

```
config ports <portlist> tunnel hdlc [off | mpls]
```

Description

Enables tunneling for HDLC encapsulated frames from a SONET port through an MPLS network.

Syntax Description

portlist	Specifies the SONET port number(s).
off	Disables HDLC tunneling.
mlps	Enables an MLPS TLS-tunnel.

Default

The default is `off`.

Usage Guidelines

The ingress SONET port encapsulates the entire HDLC frame (including the HDLC header and FCS) inside an Ethernet/MPLS header. The egress SONET port strips the Ethernet/MPLS header and forwards the HDLC frame. HDLC idle bytes (x7E) are not tunneled, but runts and aborted frames are. HDLC control bytes are destuffed on ingress and stuffed on egress.

When a SONET port is configured for HDLC tunneling, PPP should not be configured on the port (BCP and IPCP should be off). Furthermore, the port should be the only port in a VLAN and a MPLS TLS-tunnel should be configured for this VLAN. The payload inside HDLC could be PPP or some other HDLC-encapsulated protocol. SONET APS (automatic protection switching) is supported between tunneled PoS ports on the same module or different modules in the same switch. APS for tunneled ports is not supported for ports on different switches.

Example

The following command example configures an HDLC tunnel, and applies to a PoS module installed in slot 1 of a BlackDiamond switch:

```
config ports 1:4 tunnel hdlc mpls
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config ppp ports

```
config ppp [bcp [on | off] | ipcp [on {peer-ipaddress <ip address>} | off]]
ports <portlist>
```

Description

Configures the network control protocol that will run on the specified PPP ports.

Syntax Description

bcp	Specifies bridging control protocol for the port.
ipcp	Specifies IP control protocol for the port.
on	Enables the designated protocol on the port.
off	Disables the designated protocol on the port.
peer-ipaddress	Allows you to configure IP address of the peer router.
ip address	Specifies IP address of the peer router.
portlist	Specifies the port number(s).

Default

By default, BCP is enabled on all PoS ports. (However, ports 2 and 4 of OC-3c modules are not members of any VLANs by default; all other ports are members of the default VLAN by default.)

Usage Guidelines

The bcp keyword represents the bridging control protocol (BCP), and the ipcp keyword represents the IP control protocol. IPCP and BCP are mutually exclusive configuration options for a given port (that is, they cannot both be enabled simultaneously on the same port). Generally, when IPCP is enabled on a port, the port must be a member of one and only one VLAN. Furthermore, no other ports can be members of this VLAN, and IP routing is the only protocol supported on the VLAN. The one exception to this rule occurs when SONET automatic protection switching (APS) is enabled. A single VLAN can contain two IPCP-enabled ports if they are members of the same APS group.

The peer-ipaddress keyword provides an option to configure the IP address of the peer router. This can be useful with peer routers that do not advertise their IP address using the IPCP IP address configuration option (for example, Juniper routers). If the peer router does advertise an IP address via IPCP, the configured peer-ipaddress is ignored.

BCP/IPCP enables Ethernet MAC frames to be transported across a PPP link. Thus, any protocol can be transported across a BCP connection. Essentially, BCP enables the PPP link to appear as an Ethernet LAN segment to the rest of the switch. Therefore, the port may be a member of multiple VLANs, and frames can be either bridged or routed onto the link. There are restrictions regarding which ports can be bridged together (in other words, they may be members of the same VLAN) on the OC-3 PoS Module. Ports 1 and 2 on the same OC-3 module cannot be bridged together (unless they are members of the same APS group). Additionally, ports 3 and 4 on the same OC-3 module cannot be bridged together (unless they are members of the same APS group). There are no similar restrictions regarding bridging ports together on the OC-12 PoS Module.

BCP operation requires at least one Ethernet I/O module be operational in the chassis. IPCP cannot be enabled on a port unless BCP is off, and vice versa. IPCP is recommended when a PoS port only carries

routed IP traffic (because IPCP imposes less header overhead, the maximum link throughput is higher than with BCP).

Example

The following command example configures BCP on the PPP port, and applies to a PoS module installed in slot 1 of a BlackDiamond switch:

```
config ppp bcp off port 1:4
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config ppp authentication ports

```
config ppp authentication [off | chap | pap | chap-pap] ports <portlist>
```

Description

Configures authentication on the specified PPP ports.

Syntax Description

off	Disables authentication
chap	Authenticates the peer using the challenge handshake authentication protocol (CHAP).
pap	Authenticates the peer using the password authentication protocol.
chap-pap	Specifies that either CHAP or PAP may be used to authenticate the peer.
portlist	Specifies the port number(s).

Default

The default is authentication `off`.

Usage Guidelines

When `off` is specified, the peer is not authenticated. When `chap` is specified, the peer is authenticated using the challenge handshake authentication protocol (CHAP). When `pap` is specified, the peer is authenticated via the password authentication protocol (PAP). Specification of `chap-pap` indicates that either CHAP or PAP may be used to authenticate the peer.

Example

The following command example turns on CHAP authentication for port 1 of the PoS module installed in slot 8 of the BlackDiamond switch:

```
config ppp authentication chap ports 8:1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config ppp delayed-down-time ports

```
config ppp delayed-down-time <seconds> ports <portlist>
```

Description

Configures the delayed-down-time interval used by PPP for the specified ports.

Syntax Description

seconds	Specifies interval for delayed-down-time in seconds.
portlist	Specifies the port number(s).

Default

The default value is 1 second.

Usage Guidelines

The delayed-down-time interval is the amount of time that PPP waits before declaring a port down after a physical link failure has been detected. A non-zero value is useful when recovery from the link failure is fast (for example, when APS is enabled on a SONET port). In this case, APS may be able to recover from the link failure fast enough that there is no need to perturb the logical connection with the peer PPP entity, which minimizes network down time. A non-zero value is desirable when APS is configured at either end of the link. The delayed-down-time parameter is configured in seconds, with a valid range of [0..20].

Example

The following command example sets the delayed-down-time interval to 2 seconds for port 1 of the PoS module installed in slot 8 of the BlackDiamond switch:

```
config ppp delayed-down-time 2 ports 8:1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config ppp echo ports

```
config ppp echo [<seconds> <consecutive_misses> | off] ports <portlist>
```

Description

Configures the link maintenance protocol on the specified ports.

Syntax Description

seconds	Specifies the amount of time in seconds between transmissions of echo-request packets.
consecutive_ misses	Controls the amount of time that PPP waits for a reply.
off	Disables the link maintenance protocol.
portlist	Specifies the port number(s).

Default

The link maintenance protocol is `off` by default.

Usage Guidelines

When link maintenance is enabled and the port is receiving no packets, echo-request packets are transmitted over the link on a periodic basis. The `seconds` parameter in an integer in the range [1..300] that specifies the amount of time between transmissions of echo-request packets. The `consecutive_misses` parameter is an integer in the range [1..100] that controls the amount of time that PPP waits for a reply. If an echo-reply is not received within an interval of duration (`consecutive_misses * seconds`) seconds, the link is brought down. The link maintenance protocol may be disabled using the `off` keyword.

Example

The following example enables link maintenance on port 1 of a MPLS module in slot 8 and sets `seconds` to 3 and `consecutive misses` to 10:

```
config ppp echo 3 10 ports 8:1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config ppp pos checksum ports

```
config ppp pos checksum [32 | 16] ports <portlist>
```

Description

Configures the size of the HDLC Frame Check Sequence (FCS) to be used on the specified SONET ports.

Syntax Description

16 or 32	Specifies the size of the HDLC frame check sequence (either 32 bits or 16 bits).
portlist	Specifies the port number(s).

Default

The default is a 32-bit FCS.

Usage Guidelines

The two choices are a 32-bit FCS or a 16-bit FCS. RFC 2615 recommends that a 32-bit FCS be used.

Example

The following command example sets the FCS to 16 for port 1 of the PoS module installed in slot 8 of the BlackDiamond switch:

```
config ppp pos checksum 16 ports 8:1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config ppp pos scrambling ports

```
config ppp pos scrambling [on | off] ports <portlist>
```

Description

Specifies whether the payload data should be scrambled on the specified ports. RFC 2615 recommends that the SONET payload be scrambled.

Syntax Description

on	Enables scrambling.
off	Disables scrambling.
portlist	Specifies the port number(s).

Default

The default is scrambling on.

Usage Guidelines

The option of disabling scrambling is provided for backward compatibility with an earlier (now obsoleted) PoS standard specified in RFC 1619. Scrambling was introduced in RFC 2615 to alleviate potential security problems where malicious users might generate packets with bit patterns that create SONET synchronization problems.

Example

The following command example turns off the scrambling function for port 1 of the PoS module installed in slot 8 of the BlackDiamond switch:

```
config ppp pos scrambling off ports 8:1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config ppp quality ports

```
config ppp quality [off | <required_percent> {<seconds>}] ports <portlist>
```

Description

Configures the Link Quality Monitoring (LQM) protocol on the specified ports.

Syntax Description

off	Disables link quality monitoring protocol.
required_percent	Specifies required link drop percentage for link quality management (LQM).
seconds	Specifies how often (in seconds) the quality reports are to be received from the peer LQM entity.
portlist	Specifies the port number(s).

Default

The default value of `seconds` is 30. By default, LQM is `off`.

Usage Guidelines

LQM periodically transmits counts of packets/octets that were transmitted, along with counts of packets/octets that were successfully received. This information enables LQM to determine the percentage of data that is being dropped due to poor link quality. If the drop percentage is greater than (100 - `required_percent`), all network-layer protocols running over the link are brought down. You may want to bring a poor-quality link down when an alternate network path exists, or when billing is based on the amount of data transmitted. The `required_percent` parameter is an integer in the range [1..99]. The `seconds` parameter is an integer in the range [1..300] that determines how often quality reports are to be received from the peer LQM entity (that is, the reporting interval). Specifying the `seconds` parameter is optional. It can take up to seven reporting intervals for LCP to bring a link down. If the link quality subsequently improves, LCP will automatically bring the link back up; this type of service restoration will take a minimum of 7 reporting intervals.

Example

The following example enables the LQM protocol on port 1 of a MPLS module in slot 3 and sets `required_percent` to 95. Because no value is specified for the optional `seconds` parameter, the command uses the default of 30 seconds:

```
config ppp quality 95 ports 3:1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config ppp user ports

```
config ppp user <name> {encrypted} {<password>} ports <portlist>
```

Description

Configures the user name and password that the specified PPP port uses in the event the PPP peer requests authentication.

Syntax Description

name	Specifies user name for PPP peer authentication requests.
encrypted	This parameter option should not be entered.
password	Specifies the password for PPP peer authentication requests.
portlist	Specifies the port number(s).

Default

The default value of both name and password is **extreme**.

Usage Guidelines

The name is also sent when a port transmits a CHAP authentication request. The implementation responds to either CHAP or PAP authentication requests issued by the peer regardless of whether the port is configured to authenticate the peer. The name parameter is a string with a length in the range of [1..32] characters. The password parameter is also a character string, with a maximum length of 32 characters. If no password is provided on the command line, then you are prompted to enter the password twice (with the second time serving as a confirmation). You should not enter the encrypted parameter option (it is used by the switch when generating an ASCII configuration).

Example

The following command example sets the name to titus and sets the password to 1Afortune for port 1 of the PoS module installed in slot 8 of the BlackDiamond switch:

```
config ppp user "titus" "1Afortune" ports 8:1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config qosprofile

```
config qosprofile <qosprofile> {minbw <percent>} {maxbw <percent>}
{priority <level>} {minbuf <percent>} {maxbuf <percent>} {<portlist>}
{egress | ingress}
```

Description

Configures a QoS profile.

Syntax Description

qosprofile	Specifies the QoS profile to be configured.
minbw	Specifies the minimum percentage of the bandwidth available for transmissions from the profile.
maxbw	Specifies the maximum percentage of the bandwidth that can be used for transmissions from the profile.
priority	Specifies which traffic is scheduled when bandwidth is still available after the minimum requirements of all profiles have been satisfied.
level	Specifies the priority level (low, lowHi, normal, normalHi, medium, mediumHi, high, or highHi).
minbuf	This keyword is not applicable to SONET ports.
maxbuf	This keyword is not applicable to SONET ports.
portlist	Specifies the port number(s).
egress	Specifies that the flow is from the SONET port.
ingress	Specifies that the flow is to the SONET port.

Default

Normal.

Usage Guidelines

The optional `egress` and `ingress` keywords have been added to support the PoS module. These new keywords are currently only applicable to PoS ports. The PoS modules support eight egress queues and eight ingress queues per port, and the scheduling parameters for these queues are controlled by QoS profiles `qp1-qp8` (in other words, queue #0 is controlled by `qp1`, queue #1 by `qp2`, and so on). The `portlist` parameter allows QoS profiles to be customized on a SONET-port basis, while the `egress` and `ingress` keywords enable even finer customization (down to a particular `egress` or `ingress` queue on a given port). If the `egress` and `ingress` keywords are omitted, then the configured parameters apply to the `egress` queue associated with the specified `qosprofile`.

The `minbw` parameter is an integer in the range [0..100] that specifies the minimum percentage of the bandwidth that must be available for transmissions from the profile. The sum of the `minbw` parameters across all eight profiles cannot exceed 90%.

The maxbw parameter is also an integer in the range [1..100] that specifies the maximum percentage of the bandwidth that can be used for transmissions from the profile. The priority level may be set to low, lowHi, normal, normalHi, medium, mediumHi, high, or highHi. The priority determines which traffic is scheduled when bandwidth is still available after the minimum requirements of all profiles have been satisfied.

The `minbuf` and `maxbuf` keywords are not applicable to PoS ports.

Example

The following command configures the QoS profile in the egress direction, with a minimum bandwidth of 10 percent and a maximum of 20 percent:

```
config qosprofile qp8 minbw 10 maxbw 20 2:1-2:2 egress
```

History

This command was modified in an ExtremeWare IP Technology Services Release based on v6.1.5b20 to support the PoS module.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config red

```
config red [drop-probability | low-drop-probability |
high-drop-probability] <percent> {ports <portlist>}
```

Description

Configures the RED drop probability for a specified port.

Syntax Description

drop-probability	Specifies both the high and low drop probability rates.
low-drop-probability	Sets the low drop probability rate.
high-drop-probability	Sets the high drop probability rate.
percent	Specifies the percentage for the drop probability.
portlist	Specifies the port number(s).

Default

For PoS ports, both the low and high drop-probabilities default to 10%.

Usage Guidelines

The optional `low-drop-probability`, `high-drop-probability`, and `ports` keywords have been added to support the PoS module. Currently, these new keywords are only supported for SONET ports. Omission of the `ports` keyword indicates that the setting is to be applied to all ports.

The drop probability is specified as a percentage, where the `percent` parameter is an integer in the range [1..100]. The implementation provides weighted RED (WRED) functionality via support for two different drop probabilities: a `low-drop-probability` and a `high-drop-probability`. The DSCPs of IP packets indicate whether the packet should be dropped with low probability or high probability, and the appropriate percentage is then applied if WRED is active. WRED is only applied to IP packets, and the `config diffserv examination code-point` command supports complete flexibility in assigning DSCPs to the two different drop-probability levels. The configured mapping of DSCPs to drop-probability levels is used by WRED even if `diffserv examination` is disabled on the port.

The `drop-probability` keyword indicates that the specified percentage should be used for both the low and high drop-probabilities, which effectively disables WRED and reverts to standard RED operation. RED is active when the average queue length is between the minimum and maximum thresholds. In this region, the probability that a given packet is dropped increases linearly up to the configured drop probability at the maximum threshold. All packets are dropped when the average queue length exceeds the maximum threshold.

Example

The following command configure a RED high drop-probability of 20% on the SONET ports:

```
config red high-drop-probability 20 ports 2:1-2:2
```

History

This command was modified in an ExtremeWare IP Technology Services Release based on v6.1.5b20 to support PoS modules.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config red min-threshold ports

```
config red min-threshold <percent> ports <portlist>
```

Description

Configures the minimum queue length threshold for RED operation on the specified PoS ports.

Syntax Description

percent	Specifies the percentage for the minimum queue length threshold for RED operation.
portlist	Specifies the port number(s).

Default

By default, `min-threshold` is 10% for PoS ports.

Usage Guidelines

When this threshold is exceeded, the RED algorithm is activated. Currently, the command is only applicable to PoS ports. The `ports` keyword allows the threshold parameter to be configured on a PoS-port basis. The *min-threshold* is specified as a percentage, where the `percent` parameter is an integer in the range [1..100]. For PoS ports, the minimum threshold is a percentage of 1000 packet buffers, and the maximum threshold is set to minimum ((3 * minimum threshold buffers), maximum available buffers). The settings for both the minimum and maximum thresholds, in terms of number of buffers, are displayed by the `show ports info detail` command.

Example

The following command configures minimum queue length threshold of 50 for port 1 of the PoS module installed in slot 8 of the BlackDiamond switch:

```
config red min-threshold 50 port 8:1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config sonet clocking ports

```
config sonet clocking [line | internal] ports <portlist>
```

Description

Configures the clocking source for the specified SONET ports.

Syntax Description

line	Sets the line clocking on the specified port.
internal	Sets internal clocking on the specified port.
portlist	Specifies the port number(s).

Default

The default setting is `internal`.

Usage Guidelines

The clock is recovered from the received bitstream when `line` clocking is configured.

Example

The following command example selects line clocking for port 1 of the PoS module installed in slot 8 of the BlackDiamond switch:

```
config sonet clocking line ports 8:1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config sonet framing ports

```
config sonet framing [sonet | sdh] ports <portlist>
```

Description

Configures the framing type for the specified SONET ports.

Syntax Description

sonet	Sets the framing type to SONET.
sdh	Sets the framing type to SDH.
portlist	Specifies the port number(s).

Default

The default setting is `sonet`.

Usage Guidelines

You can configure each port for framing that complies with either the SONET standard or the SDH standard. SONET is primarily an American standard; SDH is the international version.

Example

The following command example selects SDH framing for port 1 of the PoS module installed in slot 8 of the BlackDiamond switch:

```
config sonet framing sdh ports 8:1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config sonet signal label ports

```
config sonet signal label [auto | <hex_octet>] ports <portlist>
```

Description

Configures the signal label value for the specified SONET ports.

Syntax Description

auto	Enables the signal label field to be automatically set.
hex_octet	Allows you to set the signal label field to a particular hex octet value.
portlist	Specifies the port number(s).

Default

The default is `auto`, where the value of the signal Label field is automatically set based on standard conventions for the given payload type.

Usage Guidelines

The signal label field occupies one byte of the path overhead associated with each SONET frame. It is used to indicate the type of contents carried in the SPE. For example, 0x16 indicates scrambled PPP/HDLC, while 0xCF indicates unscrambled PPP/HDLC. The default may be overridden by specifying a particular `hex octet` that is to be used instead, where `hex octet` is a hexadecimal integer in the range [0..xFF]. It may be necessary to specify a particular `hex octet` in order to interoperate with implementations that do not follow the standard conventions for the signal label field.

Example

The following command example sets the Signal Label to the hexadecimal value `CF` for port 1 of the PoS module installed in slot 8 of the BlackDiamond switch:

```
config sonet signal label CF ports 8:1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config sonet threshold signal degrade ports

```
config sonet threshold signal degrade <error_rate> ports <portlist>
```

Description

Configures the signal degrade threshold for the specified SONET ports.

Syntax Description

error_rate	Sets the threshold for the bit error rate for the SONET line.
portlist	Specifies the port number(s).

Default

The default is 10^{-6}.

Usage Guidelines

A signal degrade (SD) event is generated if the bit error rate (BER) for the SONET line exceeds the configured threshold. If automatic protection switching (APS) is enabled on the port, a SD event will initiate a line switch. The `error_rate` parameter is an integer in the range [5-9], where the SD BER is 10-`error_rate`. The default value of the `error_rate` parameter is 6, which equates to a SD BER of 10^{-6}, or 1 per million.

Example

The following command example sets the Signal Degrade threshold value to `8` for port 1 of the PoS module installed in slot 8 of the BlackDiamond switch:

```
config sonet threshold signal degrade 8 ports 8:1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config sonet threshold signal fail ports

```
config sonet threshold signal fail <error_rate> ports <portlist>
```

Description

Configures the signal failure threshold for the specified SONET ports.

Syntax Description

error_rate	Sets the signal failure threshold for the SONET ports.
portlist	Specifies the port number(s).

Default

The default is 10^{-5}.

Usage Guidelines

A signal failure (SF) event is generated if the bit error rate (BER) for the SONET line exceeds the configured threshold. A SF event brings the port down. If automatic protection switching (APS) is enabled on the port, a SF event will initiate a line switch. The `error_rate` parameter is an integer in the range [3-5], where the SF BER is 10^{-error_rate}. The default value of the `error_rate` parameter is 5, which equates to a SF BER of 10^{-5}, or 1 per hundred thousand.

Example

The following command example sets the signal fail threshold value to 3 for port 1 of the PoS module installed in slot 8 of the BlackDiamond switch:

```
config sonet threshold signal fail 3 ports 8:1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config sonet trace path ports

```
config sonet trace path <id_string> ports <portlist>
```

Description

Configures the path trace identifier string for the specified SONET ports.

Syntax Description

id_string	Specifies the path trace identifier string for the SONET ports.
portlist	Specifies the port number(s).

Default

The default is null.

Usage Guidelines

Path trace is a maintenance feature of SONET. One byte of the path overhead associated with each SONET frame is used to carry information identifying the originating path terminating equipment (PTE). The `id_string` parameter is a string that may contain up to 64 characters (which always includes a carriage return and a line feed character at the end). By default, `id_string` contains an IP address assigned to the VLAN that the port is a member of. This IP address is represented in dotted-decimal notation. If no IP address is assigned to the port's VLAN, `id_string` defaults to a string of 64 NULL characters. When SONET framing is configured, a 64-character string is repetitively transmitted, one character per frame. If the configured string is less than 64 characters, it is padded with NULL characters. Operation is similar when SDH framing is configured, except that the maximum string length is 15 characters. If necessary, the configured `id_string` is truncated to 15 characters.

Example

The following command example sets the path trace identifier to the string `parador` for port 1 of the PoS module installed in slot 8 of the BlackDiamond switch:

```
config sonet trace path parador ports 8:1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config sonet trace section ports

```
config sonet trace section [<id_byte> | string <id_string>]
ports <portlist>
```

Description

Configures the section trace identifier for the specified SONET ports.

Syntax Description

id_byte	Configures the ID byte section trace identifier for the specified SONET port.
id_string	Configures the ID string section trace identifier for the specified SONET port.
portlist	Specifies the port number(s).

Default

The default is 1 for SONET, null for SDH.

Usage Guidelines

Section trace is a maintenance feature of SONET. One byte of the section overhead associated with each SONET frame is used to carry information identifying the transmitting equipment. The section trace identifier has two forms: an `id_byte` and an `id_string`. The `id_byte` parameter is an integer in the range [0-255], with a default value of 1. The `id_string` parameter is a string that may contain up to 15 characters. By default, `id_string` contains 15 NULL characters. The `id_byte` parameter is only applicable when SONET framing is configured. In this case, the configured `id_byte` value is transmitted in each SONET frame. Analogously, the `id_string` parameter is only applicable when SDH framing is configured. SDH framing repetitively cycles through a 15-character string, sending one character per frame. If the configured string is less than 15 characters, it is padded with NULL characters.

Example

The following command example sets the section trace identifier to the string `1800wombat` for port 1 of the PoS module installed in slot 8 of the BlackDiamond switch:

```
config sonet trace section string 1800wombat ports 8:1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

create account pppuser

```
create account pppuser <username> {encrypted} {<password>}
```

Description

Creates a local database entry that can be used to authenticate a PPP peer.

Syntax Description

username	Specifies the user name used for authentication.
encrypted	This parameter should not be used with SONET ports.
password	Specifies the password used for authentication.

Default

N / A.

Usage Guidelines

Authentication responses include a username. When a response is received, the database is searched for an entry with the specified username. The associated password is then used to validate the authentication response. This is a new application of the existing `create account` command. The `pppuser` keyword is new. The `name` parameter is a string with a length in the range [1-32] characters. The `password` parameter is also a character string, with a maximum length of 32 characters. If no `password` is provided on the command line, then you are prompted to enter the `password` twice (with the second time serving as a confirmation). You should not enter the `encrypted` parameter option (it is used by the switch when generating an ASCII configuration).

Example

The following command example sets the authentication database name to `stretch` and sets the password to `baserunner` for port 1 of the PoS module installed in slot 8 of the BlackDiamond switch:

```
create account pppuser stretch baserunner ports 8:1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

create aps

```
create aps <group#>
```

Description

Creates an APS group with the specified group number.

Syntax Description

group#	Specifies the APS group# to which the command applies.

Default

N / A.

Usage Guidelines

You specify the group# parameter, which is an integer in the range [1-65535]. The group# is used to identify the APS group. An APS group includes one working line and one protection line. The working line and protection line can reside on the same BlackDiamond switch or two different BlackDiamond switches. The group numbers must be unique across all BlackDiamond switches that are cooperating to provide the APS function. The group numbers must also be used in a consistent manner across BlackDiamond switches. For example, if the working line is assigned to group# 1 on BlackDiamond #1, and the associated protection line resides on BlackDiamond #2, then the protection line must also be assigned to group #1 on BlackDiamond #2. The group# is used to identify the partner (in other words, working or protection) line in Ethernet messages exchanged by BlackDiamond switches that are cooperating to provide the APS function.

Example

The following command example creates APS group 1001 on the BlackDiamond switch:

```
create aps 1001
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

delete aps

```
delete aps <group#>
```

Description

Deletes the specified APS group.

Syntax Description

group#	Specifies the APS group# to which the command applies.

Default

N/A.

Usage Guidelines

You specify the group# parameter, which is an integer in the range [1-65535]. The group# identifies the APS group to delete.

Example

The following command example deletes APS group 1001:

```
delete aps 1001
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

disable aps

```
disable aps
```

Description

Disables the APS function for an entire switch.

Syntax Description

This command has no arguments or variables.

Default

APS is disabled by default.

Usage Guidelines

If APS is disabled, interfaces configured as protection lines will not carry any traffic.

Example

To disable the APS function for the entire switch, use the following command:

```
disable aps
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

disable red ports queue

```
disable red ports <portlist> queue <queue#>
```

Description

Disables RED on the specified ports.

Syntax Description

portlist	Specifies the port number(s). May be in the form 1, 2, 3-5, 2:5, 2:6-8.
queue#	Specifies the queue for which the RED function is disabled. This parameter is supported for the PoS module only.

Default

Disabled.

Usage Guidelines

The queue keyword has been added to support the PoS module. Currently, this new keyword is only applicable to PoS ports. The keyword allows the RED function to be selectively enabled on an individual queue basis. The queue# parameter is an integer in the range [0-7]. If the queue keyword is omitted, then the command applies to all egress queue numbers for the PoS port(s). RED is not supported on the ingress queues.

Example

The following command disables RED for all PHBs except the EF PHB:

```
disable red ports 2:1-2:2 queue 8
```

History

This command was modified in an ExtremeWare IP Technology Services Release based on v6.1.5b20 to support PoS modules.

Platform Availability

The general form of this command is available on the "*i*" series platforms. The optional queue parameter is available only on the PoS module on a BlackDiamond 6800 series chassis-based system.

enable aps

```
enable aps
```

Description

Enables the APS function for an entire switch.

Syntax Description

This command has no arguments or variables.

Default

APS is disabled by default.

Usage Guidelines

If APS is enabled, interfaces configured as protection lines can carry traffic.

Example

To enable the APS function for the entire switch, use the following command:

```
enable aps
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

enable red ports queue

```
enable red ports <portlist> queue <queue#>
```

Description

Enables RED on the specified PoS ports.

Syntax Description

portlist	Specifies the port number(s).
queue#	Specifies the queue for which the RED function is enabled.

Default

By default, RED is disabled.

Usage Guidelines

The queue keyword has been added to support the PoS module. Currently, this new keyword is only applicable to PoS ports. The keyword allows the RED function to be selectively enabled on an individual queue basis. The queue# parameter is an integer in the range [0-7]. If the queue keyword is omitted, then the command applies to all egress queue numbers for the PoS port(s). (RED is not supported on the ingress queues.)

Example

The following command enables RED for all PHBs except the EF PHB:

```
enable red ports 2:1-2:2
```

History

This command was modified in an ExtremeWare IP Technology Services Release based on v6.1.5b20 to support PoS modules.

Platform Availability

The general form of this command is available on the "*i*" series platforms. The optional queue parameter is available only on the PoS module on a BlackDiamond 6800 series chassis-based system.

show aps

```
show aps {<group#>} {detail}
```

Description

Displays APS group status information.

Syntax Description

group#	Specifies the APS group# to which the command applies.
detail	Displays more detailed status information for the APS groups.

Default

By default, the command shows summarized status for the APS group(s).

Usage Guidelines

The user can optionally specify a group# parameter, which is an integer in the range [1-65535]. The group# identifies a particular APS group for which status is to be shown. Alternatively, you can enter **show aps** with no parameters to obtain status for all configured APS groups. More detailed status information can be obtained for the APS group(s) by specifying the detail parameter.

Summary status includes the following information for each APS group:

- Provisioned values of all APS configuration parameters, including SONET port numbers and whether the ports are performing the working or protection line function.
- An indication of whether the line associated with each configured port is active or inactive from an APS perspective, along with a timestamp indicating when the last APS state change occurred.
- An indication of whether a signal fail (SF) or signal degrade (SD) event due to an excessive bit error rate (BER) currently exists on the line associated with each configured port, along with a timestamp indicating when the last such error occurred. (Note that the BER thresholds that cause SF and SD events may be specified as part of configuring a SONET port.)
- Counts of the number of SF and SD events initiated by each configured port due to an excessive BER.
- Count of the number of APS authentication failures (that is, a count of the number of received APS control packets that have been discarded due to authentication failures).

Detailed status includes the information reported in the summary status along with additional status and management counters. Detailed status only applies to ports performing the protection line function.

Detailed management counters reported for each protection-line port include:

- Automatic line switches initiated by working-line switch
- Automatic line switches initiated by protection-line switch
- Automatic line switches initiated by ADM
- Line switches initiated due to external commands (for example, force or manual switch command)
- Line switches completed successfully

- Hello protocol failures (this count is included as a component of the automatic line switches initiated by protection-line switch counter)
- APS mode mismatch failures (occurs when the ADM indicates that it is provisioned for the 1:n APS architecture, or when the ADM indicates that it is provisioned for unidirectional-switching mode)
- Protection switching byte failures (occurs when the received K1 byte is either inconsistent or contains an invalid request code)
- Channel mismatch failures (occurs when the channel number in the transmitted K1 byte does not match the channel number in the received K2 byte)
- Far-end protection line failures (occurs when a signal fail request code is received on the protection line)

Additional detailed status information reported for each protection-line port includes:

- Current contents of received K1 and K2 bytes
- Contents of K1 and K2 bytes that are currently being transmitted
- An indication of whether an APS mode mismatch failure is currently active
- An indication of whether a protection switching byte failure is currently active
- An indication of whether a channel mismatch failure is currently active
- An indication of whether a Far-end protection line failure is currently active

Example

The following command displays APS group status information:

```
show aps
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

show flowstats

```
show flowstats {<portlist> | export {<group#>} {detail}}
```

Description

Displays status information for the flow statistics function.

Syntax Description

portlist	Specifies the port number(s).
export	Displays status information for export groups, which are configured on a switch-wide basis.
group#	Use this optional parameter with the `export` keyword to display status information for a specific export group. If you do not specify a value for the `group#` parameter, the `export` keyword by itself displays status information for all export groups.
detail	Displays detailed export group status information.

Default

By default, the command shows summarized status.

Usage Guidelines

The `portlist` parameter can be used to specify the SONET port(s) for which status is to be shown. Alternatively, you can specify the `export` keyword to obtain status for export groups, which are configured on a switch-wide basis. Status can be obtained for a specific export group, identified by the `group#` parameter, or for all export groups by omitting the `group#` parameter. Status can be obtained for all ports by omitting both the `portlist` parameter and the `export` keyword (in other words, by simply entering **`show flowstats`** with no parameters). More detailed export group status information may be obtained by specifying the detail parameter.

Summary status for a port includes the following information:

- Values of all flow statistics configuration parameters
- Count of flow records that have been exported
- Counts of the number of packets/bytes for which flow statistics were not maintained due to insufficient resources

Summary status for an export group port includes the following information:

- Values of all configuration parameters
- State of each export destination device

Detailed status for an export group includes the information reported in the summary status along with the following additional management counters:

- Counts of flow records that have been exported to each flow-collector destination
- Counts of the number of times each flow-collector destination has been taken out of service due to health-check failures

Example

The following command displays status information for the flow statistics function:

```
show flowstats
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

show ppp

```
show ppp {<portlist>} {detail}
```

Description

Displays status information for PPP ports.

Syntax Description

portlist	Specifies the port number(s).
detail	Displayed more detailed status information for the PPP ports.

Default

By default, the command shows summarized status for the PPP port(s).

Usage Guidelines

The `portlist` parameter can be used to specify the port(s) for which status is to be shown. Alternatively, you can enter **`show ppp`** with no parameters to obtain status for all PPP ports. More detailed status information can be obtained for the PPP port(s) by specifying the detail parameter.

Summary status includes the following information for each PPP port:

- Values of all PPP configuration parameters
- Physical link status
  - operational
  - down
  - LCP state
  - IPCP/BCP state
  - EDPCP state
  - link packet and octet counters

Detailed status includes the information reported in the summary status along with the following additional status and management counters:

- Detailed link status
  - PPP link phase
- Detailed LCP status
  - LCP options negotiated (local and remote)
  - LCP packet counters
  - number of link-down events due to PPP maintenance
- Detailed authentication status
  - remote username (if applicable)
  - CHAP/PAP packet counters

- Detailed IPCP/BCP status
  - options negotiated (local and remote)
  - packet counters
- Detailed LQM status
  - statistics from last received LQR (Link Quality Report)
  - time since last received LQR
  - LQR packet counters
  - number of link-down events due to LQM

Example

The following command displays status information for the PPP ports:

```
show ppp
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

show sonet

```
show sonet {<portlist>} {detail}
```

Description

Displays SONET port status.

Syntax Description

portlist	Specifies the port number(s).
detail	Displays more detailed status information for the ports.

Default

By default, the command shows summarized status for the port(s).

Usage Guidelines

You can use the `portlist` parameter to specify which SONET port(s) you want to display the status for. You can also omit the `portlist` parameter to obtain status for all SONET ports. More detailed status information can be obtained for the port(s) by specifying the detail parameter. Summary status includes the following information for each port:

- Values of all port configuration parameters
- State of the port
- Identification of all currently active events

Example

The following command displays the SONET port status:

```
show sonet
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

unconfig aps

```
unconfig aps <group#>
```

Description

Resets the APS group configuration parameters to their default values.

Syntax Description

group#	Specifies the APS group# to which the command applies.

Default

N/A.

Usage Guidelines

You specify the group# parameter, which is an integer in the range [1-65535]. The group# identifies the APS group that the command applies to. The command does not affect the ports that have been added to the APS group. The command does cancel any outstanding lockout, force, or manual switch requests.

Example

The following command example resets the configuration parameters of APS group 1001 to their default values:

```
unconfig aps 1001
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

unconfig diffserv dscp-mapping ports

```
unconfig diffserv dscp-mapping ports <portlist>
```

Description

Resets the DSCP mapping tables for the specified PoS ports to their default values.

Syntax Description

portlist	Specifies the port number(s).

Default

N / A.

Usage Guidelines

Three DSCP mapping tables are supported per SONET port. One of the tables is used in the ingress direction and two are used for egress flows (onto the SONET link). The two egress tables are for the congested and non-congested states, as determined by the RED algorithm (in other words, the congested state is when the average queue length is greater than the minimum RED threshold). If RED is not enabled on the SONET port, then the egress congested-state mapping table is not used.

The tables are very simple. In the ingress direction, the input DSCP of a packet received from the SONET link is replaced with an output DSCP before the packet is forwarded. The replacement is straightforward; the input DSCP is used as an index into a 64-entry table that contains the output DSCPs associated with each of the input DSCP values. The operation is similar in the egress direction, with the DSCP mapping occurring before the packet is transmitted onto the SONET link(s). The mapping operation is performed after the packet has been assigned to a QoS profile. One potential use of the DSCP mapping capability is reconciliation of varying DiffServ policies at the boundary between autonomous systems (for example, at the boundary between two ISPs). The availability of different tables for the congested / non-congested states is useful for marking operations that increase the drop probability of packets during times of congestion, as discussed in the DiffServ assured forwarding (AF) RFC.

This command is currently only applicable to SONET ports.

Example

The following command resets the DSCP mapping tables for port 1, slot 8 of a BlackDiamond switch to their default values:

```
unconfig diffserv dscp-mapping port 8:1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

unconfig ppp ports

```
unconfig ppp ports <portlist>
```

Description

Resets the PPP configuration parameters for the specified ports to their default values.

Syntax Description

portlist	Specifies the port number(s).

Default

N / A.

Usage Guidelines

By default, BCP is enabled on all PoS ports. (However, ports 2 and 4 of OC-3c modules are not members of any VLANs by default; all other ports are members of the default VLAN by default.)

Example

The following command resets the PPP configuration parameters for port 1, slot 8 of a BlackDiamond switch to the default values:

```
unconfig ppp ports 8:1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

unconfig sonet ports

```
unconfig sonet ports <portlist>
```

Description

Resets the configuration parameters of the specified SONET port to their default values.

Syntax Description

portlist	Specifies the port number(s).

Default

N/A.

Usage Guidelines

The following are the SONET port default values:

clock setting	internal
Framing	sonet
signal label	auto, where the value of the signal Label field is automatically set based on standard conventions for the given payload type.
threshold signal degrade	10^{-6}
threshold signal fail	10^{-5}
trace path	null
trace section	1 for SONET, null for SDH

Example

The following command resets the configuration parameters for port 1, slot 8 of a BlackDiamond switch to the default values:

```
unconfig sonet ports 8:1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.5b20.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

19 BGP Commands

Border Gateway Protocol (BGP) is an exterior routing protocol that was developed for use in TCP/IP networks. The primary function of BGP is to allow different autonomous systems (ASs) to exchange network reachability information.

An autonomous system is a set of routers that are under a single technical administration. This set of routers uses a different routing protocol (such as OSPF) for intra-AS routing. One or more routers in the AS are configured to be border routers, exchanging information with other border routers (in different autonomous systems) on behalf of all of the intra-AS routers.

BGP can be used as an exterior gateway protocol (E-BGP), or it can be used within an AS as an interior gateway protocol (I-BGP).

BGP Attributes

The following well-known BGP attributes are supported by the switch:

- Origin – Defines the origin of the route. Possible values are IGP, EGP, and incomplete.
- AS_Path – The list of ASs that are traversed for this route.
- Next_hop – The IP address of the next hop BGP router to reach the destination listed in the NLRI field.
- Multi_Exist_Discriminator – Used to select a particular border router in another AS when multiple border routers exist.
- Local_Preference – Used to advertise this router's degree of preference to other routers within the AS.
- Atomic_aggregate – Indicates that the sending border router is used a route aggregate prefix in the route update.
- Aggregator – Identifies the BGP router AS number and IP address that performed route aggregation.
- Community – Identifies a group of destinations that share one or more common attributes.
- Cluster_ID – Specifies a 4 byte field used by a route reflector to recognize updates from other route reflectors in the same cluster.

BGP Communities

A BGP community is a group of BGP destinations that require common handling. ExtremeWare supports the following well-known BGP community attributes:

- no-export
- no-advertise
- internet

BGP Features

This section describes the following BGP features supported by ExtremeWare:

- Route Reflectors
- Route Confederations
- Route Aggregation
- IGP Synchronization
- Using the Loopback Interface
- BGP Peer Groups

clear bgp neighbor counters

```
clear bgp neighbor [<ip address> | all] counters
```

Description

Resets the BGP counters for one or all BGP neighbor sessions to zero.

Syntax Description

ip address	Specifies the IP address of a specific BGP neighbor.
all	Specifies that counters for all BGP neighbors should be reset.

Default

N/A.

Usage Guidelines

This command resets the following counters:

- In-total-msgs
- Out-total-msgs
- In-updates
- Out-updates
- Last-error
- FsmTransitions

Example

The following command resets the counters for the BGP neighbor at 10.20.30.55:

```
clear bgp neighbor 10.20.30.55 counters
```

History

This command was first available in ExtremeWare 6.2.1

This command was modified in ExtremeWare 6.2.2 to add the FsmTransitions counter.

Platform Availability

This command is available on all *"i"* series platforms.

config bgp add aggregate-address

```
config bgp add aggregate-address <ip address>/<masklength> {as-set}
{summary-only} {advertise-route-map <route-map>} {attribute-route-map
<route-map>}
```

Description

Configures a BGP aggregate route.

Syntax Description

ipadddress	Specifies an IP address.
masklength	Specifies a netmask length.
as-set	Specifies to aggregate only the path attributes of the aggregate routes.
summary-only	Specifies to send only aggregated routes to the neighbors.
advertise-route-map	Specifies the route map used to select routes for this aggregated route.
attribute-route-map	Specifies the route map used to set the attributes of the aggregated route.

Default

N / A.

Usage Guidelines

Route aggregation is the process of combining the characteristics of several routes so that they are advertised as a single route. Aggregation reduces the amount of information that a BGP speaker must store and exchange with other BGP speakers. Reducing the information that is stored and exchanged also reduces the size of the routing table.

To use BGP route aggregation, follow these steps:

1. Enable aggregation using the following command:

   ```
   enable bgp aggregation
   ```

2. Create an aggregate route using the following commands:

   ```
   config bgp add aggregate-address <ip address>/<masklength> {as-set} {summary-only}
   {advertise-route-map <route-map>} {attribute-route-map <route-map>}
   ```

Example

The following command configures a BGP aggregate route:

```
config bgp add aggregate-address 192.1.1.5/30
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp add confederation-peer sub-AS-number

```
config bgp add confederation-peer sub-AS-number <sub_as_number>
```

Description

Specifies the list of sub-ASs that belong to a confederation.

Syntax Description

sub_as_list	Specifies a list of sub-ASs.

Default

N/A.

Usage Guidelines

A maximum of 16 AS numbers can be specified.

BGP requires networks to use a fully-meshed router configuration. This requirement does not scale well, especially when BGP is used as an interior gateway protocol. One way to reduce the size of a fully-meshed AS is to divide the AS into multiple sub-autonomous systems and group them into a *routing confederation*. Within the confederation, each sub-AS must be fully-meshed. The confederation is advertised to other networks as a single AS.

Example

The following command specifies the list of sub-ASs that belong to a confederation:

```
config bgp add confederation-peer sub-as-number 65002
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp add network

```
config bgp add network <ip address>/<mask_length> {<route_map>}
```

Description

Adds a network to be originated from this router.

Syntax Description

ip address	Specifies an IP address.
mask_length	Specifies a netmask length.
route_map	Specifies a route map.

Default

N/A.

Usage Guidelines

The network must be reachable by the router.

Using the `export` command to redistribute routes complements the redistribution of routes using the `config bgp add network` command. The `config bgp add network` command adds the route to BGP only if the route is present in the routing table. The `enable bgp export` command redistributes an individual route from the routing table to BGP. If you use both commands to redistribute routes, the routes redistributed using the `network` command take precedence over routes redistributed using the `export` command.

Example

The following command adds a network to be originated from this router:

```
config bgp add network 192.1.1.14/30
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp as-number

```
config bgp as-number <as_number>
```

Description

Changes the local AS number used by BGP.

Syntax Description

as_number	Specifies a local AS number.

Default

N/A.

Usage Guidelines

BGP must be disabled before the as number can be changed.

Example

The following command changes the local AS number used by BGP:

```
config bgp as-number 65001
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp cluster-id

```
config bgp cluster-id <cluster_id>
```

Description

Appends a BGP route reflector cluster ID to the cluster list of a router.

Syntax Description

cluster_id	Specifies a 4 byte field used by a route reflector to recognize updates from other route reflectors in the same cluster.

Default

N / A.

Usage Guidelines

Used when multiple route reflectors are used within the same cluster of clients.

BGP must be disabled before configuring the cluster ID.

Example

The following command appends a BGP route reflector cluster ID to the cluster list of a route:

```
config bgp cluster-id 40000
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp confederation-id

```
config bgp confederation-id <confederation_id>
```

Description

Specifies a BGP routing confederation ID.

Syntax Description

confederation_id	Specifies a routing confederation identifier.

Default

N / A.

Usage Guidelines

BGP requires networks to use a fully-meshed router configuration. This requirement does not scale well, especially when BGP is used as an interior gateway protocol. One way to reduce the size of a fully-meshed AS is to divide the AS into multiple sub-autonomous systems and group them into a *routing confederation*. Within the confederation, each sub-AS must be fully-meshed. The confederation is advertised to other networks as a single AS.

Use a confederation ID of 0 to indicate no confederation.

Example

The following command specifies the BGP routing confederation ID as *200*:

```
config bgp confederation-id 200
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp delete aggregate-address

```
config bgp delete aggregate-address [<ip address/masklength> | all]
```

Description

Deletes one or all BGP aggregated route.

Syntax Description

ip address/masklength	Specifies an IP address and netmask length.
all	Specifies all aggregated routes.

Default

N / A.

Usage Guidelines

Route aggregation is the process of combining the characteristics of several routes so that they are advertised as a single route. Aggregation reduces the amount of information that a BGP speaker must store and exchange with other BGP speakers. Reducing the information that is stored and exchanged also reduces the size of the routing table.

Example

The following command deletes a BGP aggregate route:

```
config bgp delete aggregate-address 192.1.1.5/30
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp delete confederation-peer sub-AS-number

```
config bgp delete confederation-peer sub-AS-number <sub_as_number>
```

Description

Specifies a sub-AS that should be deleted from a confederation.

Syntax Description

sub_as_list	Specifies a list of sub-ASs.

Default

N/A.

Usage Guidelines

A maximum of 16 AS numbers can be specified.

BGP requires networks to use a fully-meshed router configuration. This requirement does not scale well, especially when BGP is used as an interior gateway protocol. One way to reduce the size of a fully-meshed AS is to divide the AS into multiple sub-autonomous systems and group them into a *routing confederation*. Within the confederation, each sub-AS must be fully-meshed. The confederation is advertised to other networks as a single AS.

Example

The following command specifies the list of sub-ASs that belong to a confederation:

```
config bgp delete confederation-peer sub-as-number 65002
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp delete network

```
config bgp delete network [all | <ip address>/<masklength>]
```

Description

Deletes a network to be originated from this router.

Syntax Description

all	Specifies all networks.
ip address	Specifies an IP address and a netmask length.

Default

N/A.

Usage Guidelines

None.

Example

The following command deletes a network to be originated from this router:

```
config bgp delete network 192.1.1.14/30
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp local-preference

```
config bgp local-preference <local_preference>
```

Description

Changes the default local preference attribute.

Syntax Description

local_preference	Specifies a value used to advertise this router's degree of preference to other routers within the AS.

Default

100.

Usage Guidelines

The range is 0 to 4294967295.

BGP selects routes based on the following precedence (from highest to lowest):

- weight
- local preference
- shortest length (shortest AS path)
- lowest origin code
- lowest MED
- route from external peer
- lowest cost to Next Hop
- lowest routerID

Example

The following command changes the default local preference attribute to *500*:

```
config bgp local-preference 500
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp med

```
config bgp med [<number> | none]
```

Description

Configures the metric to be included in the Multi-Exit-Discriminator (MED) path attribute. The MED path attribute is included in route updates sent to external peers if a value is configured.

Syntax Description

number	Specifies a multi-exist-discriminator number.
none	Specifies not to use a multi-exist-discriminator.

Default

N / A.

Usage Guidelines

BGP selects routes based on the following precedence (from highest to lowest):

- weight
- local preference
- shortest length (shortest AS path)
- lowest origin code
- lowest MED
- route from external peer
- lowest cost to Next Hop
- lowest routerID

Example

The following command configures the metric to be included in the MED path attribute:

```
config bgp med 3
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp neighbor as-path-filter

```
config bgp neighbor [<ip address> | all] as-path-filter [in | out] [none |
<access_profile>]
```

Description

Configures the AS path filter for a neighbor.

Syntax Description

ip address	Specifies an IP address.
all	Specifies all neighbors.
in	Specifies to install the filter on the input side.
out	Specifies to install the filter on the output side.
none	Specifies to remove the filter.
access_profile	Specifies an access profile.

Default

N/A.

Usage Guidelines

The filter is defined using the access-profile mechanism and can be installed on the input side and/or the output side of the router.

You can use BGP peer groups to group together up to 128 BGP neighbors. All neighbors within the peer group inherit the parameters of the BGP peer group. The following mandatory parameters are shared by all neighbors in a peer group:

- remote AS
- source-interface
- out-nlri-filter
- out-aspath-filter
- out-route-map
- send-community
- next-hop-self

Example

The following command configures the AS path filter for a neighbor based on the access profile *nosales:*

```
config bgp neighbor 192.1.1.22 as-path-filter in nosales
```

History

This command was available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp neighbor maximum-prefix

```
config bgp neighbor [<ip address> | all] maximum-prefix <number>
{{threshold <percent>} {teardown {holddown-interval <seconds>}}
{send-traps}
```

Description

Configures the maximum number of IP prefixes accepted from a BGP neighbor.

Syntax Description

ip address	Specifies an IP address.
all	Specifies all neighbors.
number	Specifies the maximum number of prefixes accepted. The range is 0 to 4294967294. A value of 0 disables prefix limit feature.
percent	Specifies the percentage of the maximum prefix (threshold) at which a warning message is printed in the log.
teardown	Specifies that the peer session is torn down when the maximum is exceeded.
seconds	Specifies the length of time before the session is re-established. If the session cannot be re-established, it is kept down until the peer is enabled. The range is 30 to 86400.
send-traps	Specifies sending “number of prefix reached threshold” and “number of prefix exceed the max-prefix limit” SNMP traps.

Default

This feature is disabled by default.

The default threshold is 75%.

By default, `teardown` is not specified.

By default, `send-traps` is not specified.

Usage Guidelines

Configure the peer group before configuring the neighbors. To configure the peer group, use the following command:

`config bgp peer-group maximum-prefix`

Example

The following command configures the maximum number of IP prefixes accepted from all neighbors to 5000, sets the threshold for warning messages to 60%, and specifies SNMP traps:

```
config bgp neighbor all maximum-prefix 5000 threshold 60 send-traps
```

History

This command was introduced in ExtremeWare 6.2.2.

config bgp neighbor next-hop-self

```
config bgp neighbor [<ip address> | all] [next-hop-self | no-next-hop-self]
```

Description

Configures the next hop address used in the updates to be the address of the BGP connection originating the update.

Syntax Description

ip address	Specifies an IP address.
all	Specifies all neighbors.
next-hop-self	Specifies that the next hop address used in the updates be the address of the BGP connection originating it.
no-next-hop-self	Specifies that the next hop address used in the updates not be the address of the BGP connection originating it.

Default

N/A.

Usage Guidelines

These settings apply to the peer group and all neighbors of the peer group.

You can use BGP peer groups to group together up to 128 BGP neighbors. All neighbors within the peer group inherit the parameters of the BGP peer group. The following mandatory parameters are shared by all neighbors in a peer group:

- remote AS
- source-interface
- out-nlri-filter
- out-aspath-filter
- out-route-map
- send-community
- next-hop-self

Example

The following command configures the next hop address used in the updates to be the address of the BGP connection originating it:

```
config bgp neighbor next-hop-self
```

History

This command was available in ExtremeWare 6.1.5.

config bgp neighbor nlri-filter

```
config bgp neighbor [<ip address> | all] nlri-filter [in | out] [none |
<access_profile>]
```

Description

Configures an NLRI filter for a neighbor.

Syntax Description

ip address	Specifies an IP address.
all	Specifies all neighbors.
in	Specifies to install the filter on the input side.
out	Specifies to install the filter on the output side.
none	Specifies to remove the filter.
access_profile	Specifies an access profile.

Default

N/A.

Usage Guidelines

The NLRI filter is defined using the access-profile mechanism and can be installed on the input side and/or the output side of the router.

You can use BGP peer groups to group together up to 128 BGP neighbors. All neighbors within the peer group inherit the parameters of the BGP peer group. The following mandatory parameters are shared by all neighbors in a peer group:

- remote AS
- source-interface
- out-nlri-filter
- out-aspath-filter
- out-route-map
- send-community
- next-hop-self

Example

The following command configures the NLRI filter for a neighbor based on the access profile *nosales:*

```
config bgp neighbor 192.1.1.22 nlri-filter in nosales
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp neighbor password

```
config bgp neighbor [all | <ip address>] password [none | {encrypted}
<password>]
```

Description

Configures a password for a neighbor.

Syntax Description

all	Specifies all neighbors.
ip address	Specifies an IP address.
none	Specifies not to use a password
encrypted	This option is for use only by the switch when generating an ASCII configuration file. Specifies that the password should be encrypted when the configuration is uploaded to a file. Should not be used through the CLI.
password	Specifies a password.

Default

N/A.

Usage Guidelines

When a password is configured, TCP MD5 authentication is enabled on the TCP connection that is established with the neighbor.

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

Modifying the following parameters automatically disables and enables the neighbors before changes take effect:

- remote-as
- timer
- source-interface
- soft-in-reset
- password

The `encrypted` option is used by the switch when generating an ASCII configuration file (using the `upload configuration` command), and parsing a switch-generated configuration file (using the `download configuration` command). Do not select the encrypted option in the CLI.

Example

The following command configures the password for a neighbor as *Extreme*:

```
config bgp neighbor 192.1.1.5 password extreme
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp neighbor peer-group

```
config bgp neighbor [all | <ip address>] peer-group [<peer group> | none]
{acquire-all}]
```

Description

Configures an existing neighbor as the member of a peer group.

Syntax Description

all	Specifies all neighbors.
ip address	Specifies an IP address.
peer group	Specifies a peer group.
none	Specifies to remove the neighbor from the peer group.
acquire-all	Specifies that parameters that should be inherited by the neighbor from the peer group.

Default

By default, remote AS (if configured for the peer group), source-interface, out-NLRI-filter, out-ASpath-filter, out-route-map, send-community and next-hop-self settings are inherited.

Usage Guidelines

If `acquire-all` is not specified, only the default parameters are inherited by the peer group.

To create a new neighbor and add it to a BGP peer group, use the following command:

`create bgp neighbor <ip address> peer-group <peer group> {multi-hop}`

The new neighbor is created as part of the peer group and inherits all of the existing parameters of the peer group. The peer group must have remote AS configured.

Example

The following command configures an existing neighbor as the member of the peer group *external*:

`config bgp neighbor 192.1.1.22 peer-group external`

History

This command was available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp neighbor route-map-filter

```
config bgp neighbor [<ip address> | all] route-map-filter [in | out] [none
| <route_map>]
```

Description

Configures a route map filter for a neighbor.

Syntax Description

ip address	Specifies an IP address.
all	Specifies all neighbors.
in	Specifies to install the filter on the input side.
out	Specifies to install the filter on the output side.
none	Specifies to remove the filter.
route_map	Specifies a route map.

Default

N / A.

Usage Guidelines

The route map filter can be installed on the input or output side of the router. The route map is used to modify or filter the NLRI information and the path attributes associated with it when exchanging updates with the neighbor.

You can use BGP peer groups to group together up to 128 BGP neighbors. All neighbors within the peer group inherit the parameters of the BGP peer group. The following mandatory parameters are shared by all neighbors in a peer group:

- remote AS
- source-interface
- out-nlri-filter
- out-aspath-filter
- out-route-map
- send-community
- next-hop-self

Example

The following command configures the route-map-filter filter for a neighbor based on the access profile *nosales:*

```
config bgp neighbor 192.1.1.22 route-map-filter in nosales
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp neighbor route-reflector-client

```
config bgp neighbor [<ip address> | all] [route-reflector-client |
no-route-reflector-client]
```

Description

Configures a BGP neighbor to be a route reflector client.

Syntax Description

ip address	Specifies an IP address.
all	Specifies all neighbors.
route-reflector-client	Specifies for the BGP neighbor to be a route reflector client.
no-route-reflector-client	Specifies for the BGP neighbor not to be a route reflector client.

Default

N / A.

Usage Guidelines

Another way to overcome the difficulties of creating a fully-meshed AS is to use *route reflectors*. Route reflectors allow a single router to serve as a central routing point for the AS or sub-AS.

Use this command to implicitly define the router to be a route reflector. The neighbor must be in the same AS as the router.

A *cluster* is formed by the route reflector and its client routers. Peer routers that are not part of the cluster must be fully meshed according to the rules of BGP.

Example

The following command configures a BGP neighbor to be a route reflector client:

```
config bgp neighbor 192.1.1.5 route-reflector-client
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp neighbor send-community

```
config bgp neighbor [<ip address> | all] [send-community |
dont-send-community]
```

Description

Configures whether the community path attribute associated with a BGP NLRI should be included in the route updates sent to the BGP neighbor.

Syntax Description

ip address	Specifies an IP address.
all	Specifies all neighbors.
send-community	Specifies to include the community path attribute.
dont-send-community	Specifies not to include the community path attribute.

Default

N / A.

Usage Guidelines

A BGP community is a group of BGP destinations that require common handling. ExtremeWare supports the following well-known BGP community attributes:

- no-export
- no-advertise
- no-export-subconfed

Example

The following command includes the community path attribute associated with a BGP NLRI in the route updates sent to all BGP neighbors:

```
config bgp neighbor all send-community
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp neighbor soft-reset

```
config bgp neighbor [<ip address> | all] soft-reset {input | output}
```

Description

Applies the current input or output routing policy to the routing information already exchanged with the neighbor.

Syntax Description

ip address	Specifies an IP address
all	Specifies all neighbors.
input	Specifies to apply the input routing policy.
output	Specifies to apply the output routing policy.

Default

N/A.

Usage Guidelines

The input/output policy is determined by the NLRI-filter, AS-path-filter, and the route map configured for the neighbor on the input and/or output side of the router. This command does not affect the switch configuration.

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

Modifying the following parameters automatically disables and enables the neighbors before changes take effect:

- remote-as
- timer
- source-interface
- soft-in-reset
- password

Example

The following command applies the current input routing policy to the routing information already exchanged with the neighbor:

```
config bgp neighbor 192.1.1.5 soft-reset input
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp neighbor source-interface

```
config bgp neighbor [<ip address> | all] source-interface [any | vlan
<name>]
```

Description

Changes the BGP source interface for TCP connections.

Syntax Description

ip address	Specifies an IP address.
all	Specifies all neighbors.
any	Specifies any source interface.
name	Specifies a VLAN name.

Default

Any.

Usage Guidelines

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

Modifying the following parameters automatically disables and enables the neighbors before changes take effect:

- remote-as
- timer
- source-interface
- soft-in-reset
- password

Example

The following command changes the BGP source interface on the VLAN *accounting*:

```
config bgp neighbor 192.1.1.5 source-interface vlan accounting
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp neighbor timer

```
config bgp neighbor [<ip address> | all] timer <keepalive> <holdtime>
```

Description

Configures the BGP neighbor timers.

Syntax Description

ip address	Specifies an IP address.
all	Specifies all neighbors.
keepalive	Specifies a BGP neighbor timer keepalive time in seconds. The range is 0 to 65535 seconds.
holdtime	Specifies a BGP neighbor timer hold time in seconds. The range is 0 to 21845 seconds.

Default

The default keepalive setting is 60. The default hold time is 180.

Usage Guidelines

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

Modifying the following parameters automatically disables and enables the neighbors before changes take effect:

- remote-as
- timer
- source-interface
- soft-in-reset
- password

Example

The following command configures the BGP neighbor timers:

```
config bgp neighbor 192.1.1.5 timer 120 360
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp neighbor weight

```
config bgp neighbor [<ip address> | all] weight <weight>
```

Description

Assigns a locally-used weight to a neighbor connection for the route selection algorithm.

Syntax Description

ip address	Specifies an IP address.
all	Specifies all neighbors.
weight	Specifies a BGP neighbor weight.

Default

0.

Usage Guidelines

All routes learned from this peer are assigned the same weight. The route with the highest weight is more preferable when multiple routes are available to the same network. The range is 0 to 4294967295.

BGP selects routes based on the following precedence (from highest to lowest):

- weight
- local preference
- shortest length (shortest AS path)
- lowest origin code
- lowest MED
- route from external peer
- lowest cost to Next Hop
- lowest routerID

Example

The following command assigns a locally used weight of 10 to a neighbor connection:

```
config bgp neighbor 192.1.1.5 weight 10
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp peer-group maximum-prefix

```
config bgp peer-group <name> maximum-prefix <number> {{threshold <percent>}
{teardown {holddown-interval <seconds>}} {send-traps}
```

Description

Configures the maximum number of IP prefixes accepted from a peer group.

Syntax Description

name	Specifies a peer group.
number	Specifies the maximum number of prefixes accepted. The range is 0 to 4294967294. A value of 0 disables prefix limit feature.
percent	Specifies the percentage of the maximum prefix (threshold) at which a warning message is printed in the log.
teardown	Specifies that the peer session is torn down when the maximum is exceeded.
seconds	Specifies the length of time before the session is re-established. If the session cannot be re-established, it is kept down until the peer is enabled. The range is 30 to 86400.
send-traps	Specifies sending “number of prefix reached threshold” and “number of prefix exceed the max-prefix limit” SNMP traps.

Default

This feature is disabled by default.

The default threshold is 75%.

By default, `teardown` is not specified.

By default, `send-traps` is not specified.

Usage Guidelines

Configure the peer group before configuring the neighbors. To configure the neighbors, use the following command:

`config bgp neighbor maximum-prefix`

Example

The following command configures the maximum number of IP prefixes accepted from the peer group “external” to 5000, sets the threshold for warning messages to 60%, and specifies SNMP traps:

```
config bgp peer-group external maximum-prefix 5000 threshold 60 send-traps
```

History

This command was introduced in ExtremeWare 6.2.2.

config bgp peer-group next-hop-self

```
config bgp peer-group <peer group> [next-hop-self | no-next-hop-self]
```

Description

Configures the next hop address used in the updates to be the address of the BGP connection originating the update.

Syntax Description

peer group	Specifies a peer group.
next-hop-self	Specifies that the next hop address used in the updates be the address of the BGP connection originating it.
no-next-hop-self	Specifies that the next hop address used in the updates not be the address of the BGP connection originating it.

Default

N/A.

Usage Guidelines

These settings apply to the peer group and all neighbors of the peer group.

You can use BGP peer groups to group together up to 128 BGP neighbors. All neighbors within the peer group inherit the parameters of the BGP peer group. The following mandatory parameters are shared by all neighbors in a peer group:

- remote AS
- source-interface
- out-nlri-filter
- out-aspath-filter
- out-route-map
- send-community
- next-hop-self

Example

The following command configures the next hop address used in the updates to be the address of the BGP connection originating it:

```
config bgp peer-group external next-hop-self
```

History

This command was available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp peer-group route-reflector-client

```
config bgp peer-group <peer group> [route-reflector-client |
no-route-reflector-client]
```

Description

Configures a peer group to be a route reflector client.

Syntax Description

peer group	Specifies a peer group.
route-reflector-client	Specifies that the peer group be a route reflector client.
no-route-reflector-client	Specifies that the peer group not be a route reflector client.

Default

N / A.

Usage Guidelines

This command implicitly defines this router to be a route reflector.

The peer group must be in the same AS of this router.

Example

The following command configures the peer group *external* as a route reflector client:

```
config bgp peer-group external route-reflector-client
```

History

This command was available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp peer-group send-community

```
config bgp peer-group <peer group> [send-community | dont-send-community]
```

Description

Configures whether communities should be sent to neighbors as part of route updates.

Syntax Description

peer group	Specifies a peer group.
send-community	Specifies that communities are sent to neighbors as part of route updates.
dont-send-community	Specifies that communities are not sent to neighbors as part of route updates.

Default

N / A.

Usage Guidelines

These settings apply to the peer group and all neighbors of the peer group.

You can use BGP peer groups to group together up to 128 BGP neighbors. All neighbors within the peer group inherit the parameters of the BGP peer group. The following mandatory parameters are shared by all neighbors in a peer group:

- remote AS
- source-interface
- out-nlri-filter
- out-aspath-filter
- out-route-map
- send-community
- next-hop-self

Example

The following command configures communities to be sent to neighbors as part of route updates:

```
config bgp peer-group external send-community
```

History

This command was available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp peer-group as-path-filter

```
config bgp peer-group <peer group> as-path-filter [in | out] [none |
<access profile>]
```

Description

Configures the AS-path filters for a peer group and all neighbors of the peer group.

Syntax Description

peer group	Specifies a peer group.
in	Specifies to install the filter on the input side.
out	Specifies to install the filter on the output side.
none	Specifies to remove the filter.
access profile	Specifies an access profile.

Default

N/A.

Usage Guidelines

You can use BGP peer groups to group together up to 128 BGP neighbors. All neighbors within the peer group inherit the parameters of the BGP peer group. The following mandatory parameters are shared by all neighbors in a peer group:

- remote AS
- source-interface
- out-nlri-filter
- out-aspath-filter
- out-route-map
- send-community
- next-hop-self

Example

The following command configures the as-path filters for the peer group *external* and its neighbors using the access profile *nosales*:

```
config bgp peer-group external as-path-filter in nosales
```

History

This command was available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp peer-group nlri-filter

```
config bgp peer-group <peer group> nlri-filter [in | out] [none | <access
profile>]
```

Description

Configures the NLRI filter for a peer group and all the neighbors of the peer group.

Syntax Description

peer group	Specifies a peer group.
in	Specifies to install the filter on the input side.
out	Specifies to install the filter on the output side.
none	Specifies to remove the filter.
access profile	Specifies an access profile.

Default

N/A.

Usage Guidelines

You can use BGP peer groups to group together up to 128 BGP neighbors. All neighbors within the peer group inherit the parameters of the BGP peer group. The following mandatory parameters are shared by all neighbors in a peer group:

- remote AS
- source-interface
- out-nlri-filter
- out-aspath-filter
- out-route-map
- send-community
- next-hop-self

Example

The following command configures the NLRI filter for the peer group *external* and its neighbors using the access profile *nosales*:

```
config bgp peer-group external nlri-filter in nosales
```

History

This command was available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp peer-group password

```
config bgp peer-group <peer group> password [<password> | none]
```

Description

Configures the password for a peer group and all neighbors of the peer group.

Syntax Description

peer group	Specifies a peer group.
password	Specifies a password.
none	Specifies no password.

Default

N / A.

Usage Guidelines

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

Modifying the following parameters automatically disables and enables the neighbors before changes take effect:

- remote-as
- timer
- source-interface
- soft-in-reset
- password

Example

The following command configures the password as *Extreme* for the peer group *external* and its neighbors:

```
config bgp peer-group external password extreme
```

History

This command was available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp peer-group remote-AS-number

```
config bgp peer-group <peer group> remote-AS-number <number>
```

Description

Configures the remote AS number for a peer group and all the neighbors of the peer group.

Syntax Description

peer group	Specifies a peer group.
number	Specifies a remote AS number.

Default

N / A.

Usage Guidelines

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

Modifying the following parameters automatically disables and enables the neighbors before changes take effect:

- remote-as
- timer
- source-interface
- soft-in-reset
- password

Example

The following command configures the remote AS number for the peer group *external* and its neighbors:

```
config bgp peer-group external remote-AS-number 65001
```

History

This command was available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp peer-group route-map-filter

```
config bgp peer-group <peer group> route-map-filter [in | out] [none |
<access profile>
```

Description

Configures the route maps for a peer group and all the neighbors of the peer group.

Syntax Description

peer group	Specifies a peer group.
in	Specifies to install the filter on the input side.
out	Specifies to install the filter on the output side.
none	Specifies to remove the filter.
access profile	Specifies an access profile.

Default

N/A.

Usage Guidelines

You can use BGP peer groups to group together up to 128 BGP neighbors. All neighbors within the peer group inherit the parameters of the BGP peer group. The following mandatory parameters are shared by all neighbors in a peer group:

- remote AS
- source-interface
- out-nlri-filter
- out-aspath-filter
- out-route-map
- send-community
- next-hop-self

Example

The following command configures the route map filter for the peer group *external* and its neighbors using the access profile *nosales*:

```
config bgp peer-group external route-map-filter in nosales
```

History

This command was available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp peer-group soft-reset

```
config bgp peer-group <peer group> soft-reset {[in | out]}
```

Description

Applies the current input/output routing policy to the neighbors in the peer group.

Syntax Description

peer group	Specifies a peer group.
in	Specifies to apply the input routing policy.
out	Specifies to apply the output routing policy.

Default

N/A.

Usage Guidelines

The input/output routing policy is determined by the NLRI-filter, AS-path-filter, and the route-map configured for the neighbors in the peer group on the input/output side of the router. This command does not affect configuration of the switch.

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

Modifying the following parameters automatically disables and enables the neighbors before changes take effect:

- remote-as
- timer
- source-interface
- soft-in-reset
- password

Example

The following command applies the current input routing policy to the neighbors in the peer group *external*:

```
config bgp peer-group external soft-reset in
```

History

This command was available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp peer-group source-interface

```
config bgp peer-group <peer group> source-interface [any | vlan]
```

Description

Configures the source interface for a peer group and all the neighbors of the peer group.

Syntax Description

peer group	Specifies a peer group.
any	Specifies any source interface.
vlan	Specifies a VLAN name.

Default

N / A.

Usage Guidelines

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

Modifying the following parameters automatically disables and enables the neighbors before changes take effect:

- remote-as
- timer
- source-interface
- soft-in-reset
- password

Example

The following command configures the source interface for the peer group *external* and its neighbors on the VLAN *accounting*:

```
config bgp peer-group external source-interface accounting
```

History

This command was available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp peer-group timer

```
config bgp peer-group <peer group> timer keep-alive <seconds> hold-time
<seconds>
```

Description

Configures the keepalive timer and hold timer values for a peer group and all the neighbors of the peer group.

Syntax Description

peer group	Specifies a peer group.
seconds	Specifies a keepalive time in seconds.
seconds	Specifies a hold-time in seconds.

Default

N/A.

Usage Guidelines

Changes made to the parameters of a peer group are applied to all neighbors in the peer group.

Modifying the following parameters automatically disables and enables the neighbors before changes take effect:

- remote-as
- timer
- source-interface
- soft-in-reset
- password

Example

The following command configures the keepalive timer and hold timer values for the peer group *external* and its neighbors:

```
config bgp peer-group external timer keep-alive 30 hold-time 30
```

History

This command was available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp peer-group weight

```
config bgp peer-group <peer group> weight <number>
```

Description

Configures the weight for the peer group and all the neighbors of the peer group.

Syntax Description

peer group	Specifies a peer group.
number	Specifies a BGP peer group weight.

Default

N/A.

Usage Guidelines

BGP selects routes based on the following precedence (from highest to lowest):

- weight
- local preference
- shortest length (shortest AS path)
- lowest origin code
- lowest MED
- route from external peer
- lowest cost to Next Hop
- lowest routerID

Example

The following command configures the weight for the peer group *external* and its neighbors:

```
config bgp peer-group external weight 5
```

History

This command was available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp router-id

```
config bgp router-id <router_id>
```

Description

Changes the router identifier.

Syntax Description

router_id	Specifies a router identifier.

Default

N / A.

Usage Guidelines

BGP must be disabled before changing the router ID.

BGP selects routes based on the following precedence (from highest to lowest):

- weight
- local preference
- shortest length (shortest AS path)
- lowest origin code
- lowest MED
- route from external peer
- lowest cost to Next Hop
- lowest router ID

Example

The following command changes the router ID:

```
config bgp router-id 192.1.1.13
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config bgp soft-reconfiguration

```
config bgp soft-reconfiguration
```

Description

Immediately applies the route map associated with the network command, aggregation, and redistribution.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command does not affect the switch configuration.

Example

The following command applies the route map associated with the network command, aggregation and redistribution:

```
config bgp soft-reconfiguration
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

create bgp neighbor peer-group

```
create bgp neighbor <ip address> peer-group <peer group> {multi-hop}
```

Description

Creates a new neighbor and makes it part of the peer group.

Syntax Description

ip address	Specifies an IP address.
peer group	Specifies a peer group.
multi-hop	Specifies to allow connections to EBGP peers that are not directly connected.

Default

N / A.

Usage Guidelines

All the parameters of the neighbor are inherited from the peer group. The peer group should have the remote AS configured.

To add an existing neighbor to a peer group, use the following command:

```
config bgp neighbor [<ip address> | all] peer-group <peer group> {acquire-all}
```

If you do not specify acquire-all, only the mandatory parameters are inherited from the peer group. If you specify acquire-all, all of the parameters of the peer group are inherited. This command disables the neighbor before adding it to the peer group.

Example

The following command creates a new neighbor and makes it part of the peer group *external*:

```
create bgp neighbor 192.1.1.22 peer-group external
```

History

This command was available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all "*i*" series platforms.

create bgp neighbor remote-as

```
create bgp neighbor <ip address> remote-as <as_number> {multihop}
```

Description

Creates a new BGP peer.

Syntax Description

ip address	Specifies an IP address.
as_number	Specifies a remote AS number.
multihop	Specifies to allow connections to EBGP peers that are not directly connected.

Default

N / A.

Usage Guidelines

If the AS number is the same as the AS number provided in the `enable bgp` command, then the peer is consider an IBGP peer, otherwise the neighbor is an EBGP peer. The BGP session to a newly created peer is not started until the `enable bgp neighbor` command is issued.

Example

The following command creates a new BGP peer:

```
create bgp neighbor 192.1.1.17 remote-as 65001
```

History

This command was available in ExtremeWare 6.2.

Platform Availability

This command is available on all "*i*" series platforms.

create bgp peer-group

```
create bgp peer-group <peer group>
```

Description

Creates a new peer group.

Syntax Description

peer group	Specifies a peer group.

Default

N / A.

Usage Guidelines

You can use BGP peer groups to group together up to 128 BGP neighbors. All neighbors within the peer group inherit the parameters of the BGP peer group. The following mandatory parameters are shared by all neighbors in a peer group:

- remote AS
- source-interface
- out-nlri-filter
- out-aspath-filter
- out-route-map
- send-community
- next-hop-self

Each BGP peer group is assigned a unique name when the peer group is created.

Example

The following command creates a new peer group named *external*:

```
create bgp peer-group external
```

History

This command was available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all "*i*" series platforms.

delete bgp neighbor

```
delete bgp neighbor [<ip address> | all]
```

Description

Deletes one or all BGP neighbors.

Syntax Description

ip address	Specifies an IP address.
all	Specifies all neighbors.

Default

N / A.

Usage Guidelines

Use this command to delete one or all BGP neighbors.

Example

The following command deletes the specified BGP neighbor:

```
delete bgp neighbor 192.1.1.17
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all "*i*" series platforms.

delete bgp peer-group

```
delete bgp peer-group <peer group>
```

Description

Deletes a peer group.

Syntax Description

peer group	Specifies a peer group.

Default

N/A.

Usage Guidelines

Use this command to delete a specific BGP peer group.

Example

The following command deletes the peer group named *external*:

```
delete bgp peer-group external
```

History

This command was available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all "*i*" series platforms.

disable bgp

```
disable bgp
```

Description

Disables BGP.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to disable BGP on the router.

Example

The following command disables BGP:

```
disable bgp
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

disable bgp aggregation

```
disable bgp aggregation
```

Description

Disables BGP route aggregation.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Route aggregation is the process of combining the characteristics of several routes so that they are advertised as a single route. Aggregation reduces the amount of information that a BGP speaker must store and exchange with other BGP speakers. Reducing the information that is stored and exchanged also reduces the size of the routing table.

Use this command to disable BGP route aggregation.

Example

The following command disables BGP route aggregation:

```
disable bgp aggregation
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

disable bgp always-compare-med

```
disable bgp always-compare-med
```

Description

Disables Multi Exit Discriminator (MED) from being used in the route selection algorithm.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

MED is only used when comparing paths from the same AS. Use this command to disable the MED from being used when selecting a route.

Example

The following command disables MED from being used in the route selection algorithm:

```
disable bgp always-compare-med
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

disable bgp community format

```
disable bgp community format <as-number:number>
```

Description

Disables the AS-number:number format of display for communities in the output of show and upload commands.

Syntax Description

as-number:number	Specifies an as-number:number.

Default

N/A.

Usage Guidelines

Using this command, communities are displayed as a single decimal value.

Example

The following command disables the AS-number:number format of display for communities:

```
disable bgp community format
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all "*i*" series platforms.

disable bgp export

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

Description

Disables BGP from exporting routes from other protocols to BGP peers.

Syntax Description

direct	Specifies direct routing.
static	Specifies static routing.
rip	Specifies RIP routing.
ospf	Specifies OSPF routing.
ospf-intra	Specifies OSPF-intra routing.
ospf-inter	Specifies OSPF-inter routing.
ospf-extern1	Specifies OSPF-extern1 routing.
ospf-extern2	Specifies OSPF-extern2 routing.
vip	Specifies VIP routing.

Default

N / A.

Usage Guidelines

The exporting of routes between any two routing protocols is a discreet configuration function. For example, you must configure the switch to export routes from OSPF to BGP and, if desired, you must configure the switch to export routes from BGP to OSPF. 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 BGP, and the routes to export from BGP to OSPF.

You can use route maps to associate BGP attributes including Community, NextHop, MED, Origin, and Local Preference with the routes. Route maps can also be used to filter out exported routes.

Using the `export` command to redistribute routes complements the redistribution of routes using the `config bgp add network` command. The `config bgp add network` command adds the route to BGP only if the route is present in the routing table. The `enable bgp export` command redistributes an individual route from the routing table to BGP. If you use both commands to redistribute routes, the routes redistributed using the `network` command take precedence over routes redistributed using the `export` command.

Example

The following command disables BGP from exporting routes from the OSPF protocol to BGP peers:

```
disable bgp export ospf
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

disable bgp neighbor

```
disable bgp neighbor [<ip address> | all]
```

Description

Disables the BGP session.

Syntax Description

ip address	Specifies an IP address.
all	Specifies all neighbors.

Default

N / A.

Usage Guidelines

After the session has been disabled, all the information in the route information base (RIB) for the neighbor is flushed.

Example

The following command disables the BGP session:

```
disable bgp neighbor 192.1.1.17
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

disable bgp neighbor remove-private-as-numbers

```
disable bgp neighbor [<ip address> | all] remove-private-as-numbers
```

Description

Disables the removal of private AS numbers from the AS path in route updates sent to EBGP peers.

Syntax Description

ip address	Specifies an IP address.
all	Specifies all neighbors.

Default

N / A.

Usage Guidelines

Private AS numbers are AS numbers in the range 64512 through 65534. You can remove private AS numbers from the AS path attribute in updates that are sent to external BGP (EBGP) neighbors. Possible reasons for using private AS numbers include:

- The remote AS does not have officially allocated AS numbers.
- You want to conserve AS numbers if you are multi-homed to the local AS.

Private AS numbers should not be advertised on the Internet. Private AS numbers can only be used locally within an administrative domain. Therefore, when routes are advertised out to the Internet, the routes can be stripped out from the AS paths of the advertised routes using this feature.

Example

The following command disables the removal of private AS numbers from the AS path in route updates sent to the EBGP peers:

```
disable bgp neighbor 192.1.1.17 remove-private-as-numbers
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all "*i*" series platforms.

disable bgp neighbor soft-in-reset

```
disable bgp neighbor [all | <ip address>] soft-in-reset
```

Description

Disables the soft recognition feature.

Syntax Description

all	Specifies all neighbors.
ip address	Specifies an IP address.

Default

N / A.

Usage Guidelines

Disabling the soft recognition feature can potentially limit the amount of system memory consumed by the RIB-in.

Example

The following command disables the soft recognition feature:

```
disable bgp neighbor 192.1.1.17 soft-in-reset
```

History

This command was available in ExtremeWare 6.2.

Platform Availability

This command is available on all "*i*" series platforms.

disable bgp peer-group

```
disable bgp peer-group <peer group>
```

Description

Disables a BGP peer group and all the neighbors of the peer group.

Syntax Description

peer group	Specifies a peer group.

Default

N / A.

Usage Guidelines

You can use BGP peer groups to group together up to 128 BGP neighbors. All neighbors within the peer group inherit the parameters of the BGP peer group. The following mandatory parameters are shared by all neighbors in a peer group:

- remote AS
- source-interface
- out-nlri-filter
- out-aspath-filter
- out-route-map
- send-community
- next-hop-self

Example

The following command disables the BGP peer group *external* and all of its neighbors:

```
disable bgp peer-group external
```

History

This command was available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all "*i*" series platforms.

disable bgp synchronization

```
disable bgp synchronization
```

Description

Disables the synchronization between BGP and IGP.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

When enabled, BGP waits for IGP to provide the next-hop reachability before advertising the route to an external neighbor.

Example

The following command disables the synchronization between BGP and IGP:

```
disable bgp synchronization
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all "*i*" series platforms.

disable peer-group

```
disable peer-group {soft-in-reset}
```

Description

Disables the soft recognition feature of a peer group and all the neighbors of a peer group.

Syntax Description

soft-in-reset	Specifies the soft recognition feature.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables the soft recognition feature of a peer group and its neighbors:

```
disable peer-group soft-in-reset
```

History

This command was available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all "*i*" series platforms.

enable bgp

```
enable bgp
```

Description

Enables BGP.

Syntax Description

This command has no arguments or variables.

Default

Not enabled.

Usage Guidelines

This command enables the Border Gateway Protocol (BGP) on the router.

Example

The following command enables BGP:

```
enable bgp
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

enable bgp aggregation

```
enable bgp aggregation
```

Description

Enables BGP route aggregation filtering.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Route aggregation is the process of combining the characteristics of several routes so that they are advertised as a single route. Aggregation reduces the amount of information that a BGP speaker must store and exchange with other BGP speakers. Reducing the information that is stored and exchanged also reduces the size of the routing table.

To use BGP route aggregation, follow these steps:

1. Enable aggregation using the following command:

   ```
   enable bgp aggregation
   ```

2. Create an aggregate route using the following command:

   ```
   config bgp add aggregate-address <ip address>/<masklength> {as-set} {summary-only}
   {advertise-route-map <route-map>} {attribute-route-map <route-map>}
   ```

Example

The following command enables BGP route aggregation:

```
enable bgp aggregation
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

enable bgp always-compare-med

```
enable bgp always-compare-med
```

Description

Enables BGP to use the Multi Exit Discriminator (MED) from neighbors in different autonomous systems (ASs) in the route selection algorithm.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

MED is only used when comparing paths from the same AS.

Example

The following command enables BGP to use the Multi Exit Discriminator (MED) from neighbors in different autonomous systems in the route selection algorithm:

```
enable bgp always-compare-med
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

enable bgp community format

```
enable bgp community format <as-number:number>
```

Description

Enables the as-number:number format of display for the communities in the output of `show` and `upload` commands.

Syntax Description

as-number:number	Specifies an as-number:number

Default

N/A.

Usage Guidelines

If not enabled, the communities are displayed as a single decimal value.

Example

The following command enables the AS-number:number format of display for communities:

```
enable bgp community format
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all "*i*" series platforms.

enable bgp export

```
enable bgp export [direct | static | rip | ospf | ospf-intra | ospf-inter |
ospf-extern1 | ospf-extern2 | vip] {<route map>}
```

Description

Enables BGP to export routes from other protocols to BGP peers.

Syntax Description

direct	Specifies direct routing.
static	Specifies static routing.
rip	Specifies RIP routing.
ospf	Specifies OSPF routing.
ospf-intra	Specifies OSPF-intra routing.
ospf-inter	Specifies OSPF-inter routing.
ospf-extern1	Specifies OSPF-extern1 routing.
ospf-extern2	Specifies OSPF-extern2 routing.
vip	Specifies VIP routing.
route map	Specifies a route map.

Default

Disabled.

Usage Guidelines

The exporting of routes between any two routing protocols is a discreet configuration function. For example, you must configure the switch to export routes from OSPF to BGP and, if desired, you must configure the switch to export routes from BGP to OSPF. 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 BGP, and the routes to export from BGP to OSPF.

You can use route maps to associate BGP attributes including Community, NextHop, MED, Origin, and Local Preference with the routes. Route maps can also be used to filter out exported routes.

Using the `export` command to redistribute routes complements the redistribution of routes using the `config bgp add network` command. The `config bgp add network` command adds the route to BGP only if the route is present in the routing table. The `enable bgp export` command redistributes an individual route from the routing table to BGP. If you use both commands to redistribute routes, the routes redistributed using the `network` command take precedence over routes redistributed using the `export` command.

Example

The following command enables BGP to export routes from the OSPF protocol to BGP peers:

```
enable bgp export ospf
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

enable bgp neighbor

```
enable bgp neighbor [<ip address> | all]
```

Description

Enables the BGP session. The neighbor must be created before the BGP session can be enabled.

Syntax Description

ip address	Specifies an IP address.
all	Specifies all neighbors.

Default

N / A.

Usage Guidelines

To create a new neighbor and add it to a BGP peer group, use the following command:

```
create bgp neighbor <ip address> peer-group <peer group> {multi-hop}
```

The new neighbor is created as part of the peer group and inherits all of the existing parameters of the peer group. The peer group must have remote AS configured.

To add an existing neighbor to a peer group, use the following command:

```
config bgp neighbor [<ip address> | all] peer-group <peer group> {acquire-all}
```

If you do not specify acquire-all, only the mandatory parameters are inherited from the peer group. If you specify acquire-all, all of the parameters of the peer group are inherited. This command disables the neighbor before adding it to the peer group.

Example

The following command enables the BGP session:

```
enable bgp neighbor 192.1.1.17
```

History

This command was available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

enable bgp neighbor remove-private-as-number

```
enable bgp neighbor [<ip address> | all] remove-private-as-numbers
```

Description

Enables the removal of private AS numbers from the AS path in route updates sent to EBGP peers.

Syntax Description

ip address	Specifies an IP address.
all	Specifies all neighbors.

Default

N/A.

Usage Guidelines

Private AS numbers are AS numbers in the range 64512 through 65534. You can remove private AS numbers from the AS path attribute in updates that are sent to external BGP (EBGP) neighbors. Possible reasons for using private AS numbers include:

- The remote AS does not have officially allocated AS numbers.
- You want to conserve AS numbers if you are multi-homed to the local AS.

Private AS numbers should not be advertised on the Internet. Private AS numbers can only be used locally within an administrative domain. Therefore, when routes are advertised out to the Internet, the routes can be stripped out from the AS paths of the advertised routes using this feature.

Example

The following command enables the removal of private AS numbers from the AS path in route updates sent to the EBGP peers:

```
enable bgp neighbor 192.1.1.17 remove-private-as-numbers
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all "*i*" series platforms.

enable bgp neighbor soft-in-reset

```
enable bgp neighbor [all | <ip address>] soft-in-reset
```

Description

Enables the soft recognition feature.

Syntax Description

all	Specifies all neighbors.
ip address	Specifies an IP address.

Default

N / A.

Usage Guidelines

Disabling the soft reconfiguration feature can potentially limit the amount of system memory consumed by the RIB-in.

Example

The following command enables the soft recognition feature:

```
enable bgp neighbor 192.1.1.17 soft-in-reset
```

History

This command was available in ExtremeWare 6.2.

Platform Availability

This command is available on all "*i*" series platforms.

enable bgp peer-group

```
enable bgp peer-group <peer group>
```

Description

Enables a new peer group and all the neighbors of a peer group.

Syntax Description

peer group	Specifies a peer group.

Default

N / A.

Usage Guidelines

You can use BGP peer groups to group together up to 128 BGP neighbors. All neighbors within the peer group inherit the parameters of the BGP peer group. The following mandatory parameters are shared by all neighbors in a peer group:

- remote AS
- source-interface
- out-nlri-filter
- out-aspath-filter
- out-route-map
- send-community
- next-hop-self

Example

The following command enables the BGP peer group *external* and all its neighbors:

```
enable bgp peer-group external
```

History

This command was available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all "*i*" series platforms.

enable bgp peer-group soft-in-reset

```
enable bgp peer-group <peer group> soft-in-reset
```

Description

Enables the soft recognition feature on the peer group and all the neighbors of the peer group.

Syntax Description

peer group	Specifies a peer group.

Default

N/A.

Usage Guidelines

An automatic disable and enable of the neighbor is performed if it is enabled.

Example

The following command enables the soft recognition feature on the peer group *external* and all its neighbors:

```
enable bgp peer-group external soft-in-reset
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all "*i*" series platforms.

enable bgp synchronization

```
enable bgp synchronization
```

Description

Enables synchronization between BGP and IGP.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

When enabled, BGP waits for IGP to provide the next-hop reachability before advertising the route to an external neighbor.

Example

The following command enables synchronization between BGP and IGP:

```
enable bgp synchronization
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all "*i*" series platforms.

enable peer-group

```
enable peer-group {soft-in-reset}
```

Description

Enables the soft reconfiguration feature on the peer group and all the neighbors of the peer group.

Syntax Description

soft-in-reset	Specifies the soft recognition feature.

Default

N/A.

Usage Guidelines

None.

Example

The following command enables the soft reconfiguration feature on the peer group and all its neighbors:

```
enable peer-group soft-in-reset
```

History

This command was available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all "*i*" series platforms.

show bgp

```
show bgp
```

Description

Displays BGP configuration information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Displays information such as AS number, router ID, local preference, sync flag, route reflection, cluster ID, confederation ID, and AS redistributed networks.

Example

The following command displays BGP configuration information:

```
show bgp
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

show bgp neighbor

```
show bgp neighbor <ip address> {[accepted-routes | received-routes |
rejected-routes | transmitted-routes] {detail}
[community [access-profile <access_profile> | no-advertise | no-export |
no-export-subconfed | number <community_number> | <as_id>] |
as-path [<as-path-expression> | access-profile <access_profile>] |
route-map <route map> | network <ip address>/<mask> {exact} | all]}
```

Description

Displays information about a specified neighbor.

Syntax Description

ip address	Specifies an IP address that identifies a BGP neighbor.
accepted-routes	Specifies that only accepted routes should be displayed.
received-routes	Specifies that only received routes should be displayed.
rejected-routes	Specifies that only rejected routes should be displayed.
transmitted-routes	Specifies that only transmitted routes should be displayed.
detail	Specifies to display the information in detailed format.
access_profile	Specifies an access profile used as a community attribute.
no-advertise	Specifies the no-advertise community attribute.
no-export	Specifies the no-export community attribute.
no-export-subconfed	Specifies the no-export-subconfed community attribute.
community_number	Specifies a community number.
as_id	Specifies an autonomous system ID (0-65535).
access_profile	Specifies an access profile.
route map	Specifies a route map.
ip address	Specifies an IP address.
mask	Specifies a subnet mask (number of bits).
exact	Specifies an exact match with the IP address and subnet mask.
all	Specifies all routes.

Default

N/A.

Usage Guidelines

Use this command to display information about a specific BGP neighbor. If you do not specify a neighbor, information about all neighbors is displayed.

Example

The following command displays information about a specified neighbor:

```
show bgp neighbor 10.10.10.10
```

Following is the output from this command:

```
IBGP Peer: 10.10.10.10  As: 14490  Enabled: Yes  Router: Enabled  Weight: 1
ConnectRetry: 120   HoldTimeCfg: 180   KeepaliveCfg: 60   MinAsOrig:15
Source Interface: Not configured  RRClient: No  EBGP-Multihop: No
NextHopSelf: Enabled Send Communities: No Soft Input Reconfiguration: Disabled
Max-Prefix: 100000 Threshold: 75 Teardown: Yes(HoldInt: 300) SendTraps: No
Remove Private AS : No
IN NLRI Filter     : None
OUT NLRI Filter    : None
IN AS-Path Filter  : None
OUT AS-Path Filter : None
IN ROUTE-MAP       : None
OUT ROUTE-MAP      : None
State: IDLE(Reached maximum prefix limit)
RemoteAddr:10.10.10.10:179 LocalAddr:10.10.10.51:1024 PeerRtrId:0.0.0.0
InUpdates: 26549  OutUpdates(InQ): 0(0)  InTotalMsgs: 26559  OutTotalMsgs: 9
InUpdateElapsedTime: 0:0:00:20  InMsgElapsedTime: 0:0:00:20  InPrefix: 0
HoldTime: 180  KeepAlive: 60  FsmTransitions: 1  RestartAfter: 0:04:43
FSM Down since: Mon Apr 1 15:59:42 2002 (Duration: 0:0:00:17)
LastErr: 0/0
```

History

This command was available in ExtremeWare 6.1.

This command was modified in ExtremeWare 6.2.2 to include information about maximum prefix settings.

Platform Availability

This command is available on all "*i*" series platforms.

show bgp peer-group

```
show bgp peer-group {detail | <peer group> {detail}}
```

Description

Displays the peer groups configured in the system.

Syntax Description

detail	Specifies to display the information in detailed format.
peer group	Specifies a peer group.
detail	Specifies to display the information in detailed format.

Default

N / A.

Usage Guidelines

If the `detail` keyword is specified then the parameters of the neighbors in the peer group, which are different from the ones that are configured in the peer group, will be displayed.

Example

The following command displays the peer groups configured in the system:

```
show bgp peer-group detail
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all "*i*" series platforms.

show bgp routes

```
show bgp routes {detail}
[community [access-profile <access_profile> | no-advertise | no-export |
no-export-subconfed | number <community_number> | <as_id>] |
as-path [<as-path-expression> | access-profile <access_profile>] |
route-map <route map> | network <ip address>/<mask> {exact} | all]
```

Description

Displays the BGP route information base (RIB).

Syntax Description

detail	Specifies to display the information in detailed format.
access_profile	Specifies an access profile used as a community attribute.
no-advertise	Specifies the no-advertise community attribute.
no-export	Specifies the no-export community attribute.
no-export-subconfed	Specifies the no-export-subconfed community attribute.
community_number	Specifies a community number.
as_id	Specifies an autonomous system ID (0-65535).
access_profile	Specifies an access profile.
route map	Specifies a route map.
ip address	Specifies an IP address.
mask	Specifies a subnet mask (number of bits).
exact	Specifies an exact match with the IP address and subnet mask.
all	Specifies all routes.

Default

N / A.

Usage Guidelines

None.

Example

The following command displays the BGP route information base (RIB):

```
show bgp routes all
```

History

This command was available in ExtremeWare 6.1.

Platform Availability

This command is available on all *"i"* series platforms.

20 IP Multicast Commands

IP multicast routing is a function that allows a single IP host to send a packet to a group of IP hosts. This group of hosts can include devices that reside on the local network, within a private network, or outside of the local network.

IP multicast routing consists of the following functions:

- A router that can forward IP multicast packets
- A router-to-router multicast protocol [for example, Distance Vector Multicast Routing Protocol (DVMRP) or Protocol Independent Multicast (PIM)]
- A method for the IP host to communicate its multicast group membership to a router [for example, Internet Group Management Protocol (IGMP)]

NOTE

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

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.

Distance Vector Multicast Routing Protocol (DVMRP) is a distance-vector protocol that is used to exchange routing and multicast information between routers. Like RIP, DVMRP periodically sends the entire routing table to its neighbors.

DVMRP has a mechanism that allows it to prune and graft multicast trees to reduce the bandwidth consumed by IP multicast traffic.

Protocol Independent Multicast (PIM) is a multicast routing protocol. The switch supports dense mode and sparse mode operation. You can configure dense mode or sparse mode on a per-interface basis. After they are enabled, some interfaces can run dense mode, while others run sparse mode.

PIM Dense Mode

Protocol Independent Multicast-Dense Mode (PIM-DM) is a multicast routing protocol that is similar to DVMRP. PIM-DM routers perform reverse path multicasting (RPM). However, instead of exchanging its own unicast route tables for the RPM algorithm, PIM-DM uses the existing unicast route table for the reverse path. As a result, PIM-DM requires less system memory.

PIM-DM is a broadcast and prune protocol. Using PIM-DM, multicast routes are pruned and grafted in the same way as DVMRP.

PIM Sparse Mode (PIM-SM)

Unlike PIM-DM, PIM-SM is an explicit join and prune protocol, and it supports shared trees as well as shortest path trees (SPTs). The routers must explicitly join the group(s) in which they are interested in becoming a member, which is beneficial for large networks that have group members who are sparsely distributed.

Using PIM-SM, the router sends a join message to the rendezvous point (RP). The RP is a central multicast router that is responsible for receiving and distributing multicast packets. By default, the RP is selected dynamically. You can also define a static RP in your network, using the following command:

```
config pim crp static <rp_address>
```

If you use a static RP, all switches in your network must be configured with the same RP address.

When a router has a multicast packet to distribute, it encapsulates the packet in a unicast message and sends it to the RP. The RP decapsulates the multicast packet and distributes it among all member routers.

When a router determines that the multicast rate from of a particular group from a particular originating router (not the RP) has exceeded a configured threshold, that router can send an explicit join to the originating router. When this occurs, the receiving router gets the multicast directly from the sending router, and bypasses the RP.

NOTE

You can run either PIM-DM or PIM-SM per VLAN.

PIM Mode Interoperation

An Extreme Networks switch can function as a PIM multicast border router (PMBR). A PMBR integrates PIM-SM and PIM-DM traffic.

When forwarding PIM-DM traffic into a PIM-SM network, the PMBR notifies the RP that the PIM-DM network exists. The PMBR forwards PIM-DM multicast packets to the RP, which in turn forwards the packets to those routers that have joined the multicast group.

The PMBR also forwards PIM-SM traffic to a PIM-DM network. The PMBR sends a join message to the RP and the PMBR floods traffic from the RP into the PIM-DM network.

No commands are needed to enable PIM mode interoperation. PIM mode translation is automatically enabled when a dense mode interface and a sparse mode interface are enabled on the same switch.

clear igmp snooping

```
clear igmp snooping {vlan <name>}
```

Description

Removes one or all IGMP snooping entries.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

This command can be used by network operations to manually remove IGMP snooping entries instantly. However, removing an IGMP snooping entry can disrupt the normal forwarding of multicast traffic.

Example

The following command clears IGMP snooping from VLAN *accounting*:

```
clear igmp snooping accounting
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

clear ipmc cache

```
clear ipmc cache {<group> {<source> <netmask>}}
```

Description

Resets the IP multicast cache table.

Syntax Description

group	Specifies a group address.
source	Specifies a source IP address.
netmask	Specifies a subnet mask.

Default

If no options are specified, all IP multicast cache entries are flushed.

Usage Guidelines

None.

Example

The following command resets the IP multicast table for group *224.1.2.3*:

```
clear ipmc cache 224.1.2.3
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

clear ipmc fdb

```
clear ipmc fdb [group <ipaddress> [sender <ipaddress/mask>]]
```

Description

Resets the IP multicast forwarding database entry.

Syntax Description

ipaddress	Specifies a group IP address.
ipaddress/mask	Specifies a sender IP address and netmask.

Default

Disabled.

Usage Guidelines

If no options are specified, all IP multicast forwarding database entries are cleared.

Example

The following command resets the IP multicast forwarding database entry:

```
clear ipmc fdb group 10.0.0.1 sender 10.0.0.2/24
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all platforms.

config dvmrp add

```
config dvmrp add vlan [<name> | all]
```

Description

Enables DVMRP on one or all IP interfaces.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

Disabled.

Usage Guidelines

If no VLAN is specified, DVMRP is enabled on all IP interfaces. When an IP interface is created, DVMRP is disabled by default.

Example

The following command enables DVMRP on the VLAN *accounting*:

```
config dvmrp add vlan accounting
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config dvmrp delete vlan

```
config dvmrp delete vlan [<name> | all]
```

Description

Disables DVMRP on one or all IP interfaces.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

Disabled.

Usage Guidelines

If no VLAN is specified, DVMRP is disabled on all IP interfaces.

Example

The following command disables DVMRP on the VLAN *accounting:*

```
config dvmrp delete vlan accounting
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config dvmrp timer

```
config dvmrp timer <route_report_interval> <route_replacement_time>
```

Description

Configures the global DVMRP timers.

Syntax Description

route_report_interval	Specifies the time in seconds between transmission of periodic report packets.
route_replacement_time	Specifies a hold-down time in seconds before a new route is learned.

Default

- route_report_interval default — 60 seconds.
- route_replacement_time default — 140 seconds.

Usage Guidelines

Specify the following:

- route_report_interval — The amount of time the system waits between transmitting periodic route report packets. The range is 1 to 2,147,483,647 seconds (68 years). The default setting is 60 seconds. Because triggered update is always enabled, the route report will always be transmitted prior to the expiration of the route report interval.
- route_replacement_time — The route expiration time, commonly called route timeout. Initially it is 2 x route_report_interval +20 (2 x 60 + 20 = 140). It is the time for a particular DVMRP route to expire, while the route hold-down time is initially 2 x route_report_interval (2 x 60 = 120) which is the time before a route gets removed from advertisement after it has been expired. The range is 1 to 2,147,483,647 seconds (68 years).

Example

The following command configures the DVMRP timers:

```
config dvmrp timer 300 300
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config dvmrp vlan cost

```
config dvmrp vlan [<name> | all] cost <number>
```

Description

Configures the cost (metric) of the interface.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.
number	Specifies a cost metric.

Default

The default setting is 1.

Usage Guidelines

None.

Example

The following command configures the cost (metric) of the interface on the VLAN accounting:

```
config dvmrp vlan accounting cost 5
```

History

This command was available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

config dvmrp vlan export-filter

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

Description

Configures DVMRP to filter out certain routes when performing the route advertisement.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.
access_profile	Specifies an access profile name.
none	Specifies to check the access profile for permit and deny attributes.

Default

N / A.

Usage Guidelines

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. When 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.

Example

The following command configures DVMRP to filter out certain routes according to the *nosales* access profile:

```
config dvmrp vlan accounting export-filter nosales
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

config dvmrp vlan import-filter

```
config dvmrp vlan [<name> | all] import-filter [<access_profile> | none]
```

Description

Configures DVMRP to filter certain routes received from its neighbor, and uses an access profile to determine which DVMRP routes are accepted as valid routes.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.
access_profile	Specifies an access profile name.
none	Specifies to check the access profile for permit and deny attributes.

Default

N / A.

Usage Guidelines

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. When 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.

Example

The following command configures DVMRP to filter certain routes received from its neighbor, and uses the *nosales* access profile to determine which DVMRP routes are accepted as valid routes:

```
config dvmrp vlan accounting import-filter nosales
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

config dvmrp vlan trusted-gateway

```
config dvmrp vlan [<name> | all] trusted-gateway [<access_profile> | none]
```

Description

Configures DVMRP to use the access policy to determine which DVMRP neighbor is trusted and to receive routes from.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.
access_profile	Specifies an access profile name.
none	Specifies to check the access profile for permit and deny attributes.

Default

N / A.

Usage Guidelines

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. When 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.

Example

The following command configures DVMRP to use the *nosales* access policy to determine which DVMRP neighbor is trusted and to receive routes from:

```
config dvmrp vlan accounting trusted-gateway nosales
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

config dvmrp vlan timer

```
config dvmrp vlan <name> timer <probe_interval> <neighbor_timeout_interval>
```

Description

Configures DVMRP interface timers.

Syntax Description

name	Specifies a VLAN name.
probe_interval	Specifies the time in seconds between probe messages.
neighbor_timeout_interval	Specifies the time in seconds before a neighbor router is declared to be down.

Default

The probe_interval default setting is 10 seconds. The neighbor_timeout_interval default setting is 35 seconds.

Usage Guidelines

Specify the following:

- probe_interval — The amount of time that the system waits between transmitting DVMRP probe messages. The range is 1 to 2,147,483,647 seconds (68 years). The default setting is 10 seconds.
- neighbor_timeout_interval — The amount of time before a DVMRP neighbor router is declared to be down. The range is 1 to 2,147,483,647 seconds (68 years). The default setting is 35 seconds.

Example

The following command configures the DVMRP timers:

```
config dvmrp vlan accounting timer 3000 3000
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all platforms.

config igmp

```
config igmp <query_interval> <query_response_interval>
<last_member_query_interval>
```

Description

Configures the Internet Group Management Protocol (IGMP) timers.

Syntax Description

query_interval	Specifies the interval (in seconds) between general queries.
query_response_interval	Specifies the maximum query response time (in seconds).
last_member_query_interval	Specifies the maximum group-specific query response time (in seconds).

Default

- query interval — 125 seconds
- query response interval — 10 seconds
- last member query interval — 1 second

Usage Guidelines

Timers are based on RFC2236. Specify the following:

- query_interval — The amount of time, in seconds, the system waits between sending out general queries. The range is 1 to 429,496,729 seconds.
- query_response_interval — The maximum response time inserted into the periodic general queries. The range is 1 to 25 seconds.
- last_member_query_interval — 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.

Example

The following command configures the IGMP timers:

```
config igmp 100 5 1
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config igmp snooping flood-list

```
config igmp snooping flood-list [<access profile> | none]
```

Description

Configures certain multicast addresses to be slow path flooded within the VLAN.

Syntax Description

access profile	Specifies an access profile with a list of multicast addresses to be handled. The access profile must be type ipaddress.
none	Specifies no access profile is to be used.

Default

None.

Usage Guidelines

With this command, a user can configure certain multicast addresses to be slow path flooded within the VLAN, which otherwise will be fast path forwarded according IGMP and/or layer 3 multicast protocol.

The specified access profile `<access profile>` should contain a list of addresses which will determine if certain multicast streams are to be treated specially. Typically, if the switch receives a stream with destination address which is in the `<access profile>` in 'permit' mode, that stream will be software flooded and no hardware entry would be installed.

The specified access profile must be type `ipaddess`.

When adding an IP address into the access-profile, a 32-bit host address is recommended.

This feature is meant to solve the multicast connectivity problem for unknown destination addresses within system reserved ranges. Specifically this feature was introduced to solve the problem of recognizing certain stream as control packets.

NOTE

The switch will not validate any IP address in the access profile used in this command. Therefore, slow-path flooding should be used only for streams which are very infrequent, such as control packets. It should not be used for multicast data packets. This option overrides any default mechanism of hardware forwarding (with respect to IGMP, PIM, or DVMRP) so it should be used with caution.

Slow path flooding will be done within the L2 VLAN only.

Use the none option to effectively disable slow path flooding.

You can use the `show ipconfig` command to see the configuration of slow path flooding. It will be listed in the IGMP snooping section of the display.

Example

Given access profile *access1* created as follows:

```
create access-profile access1 type ipaddress
config access-profile access1 add ipaddress 224.1.0.1/32
```

The following command configures the multicast data stream specified in *access1* for slow path flooding:

```
config igmp snooping flood-list access1
```

The following command specifies that no access profile is to be used, this effectively disabling slow path flooding:

```
config igmp snooping flood-list none
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on “i” series platforms.

config igmp snooping leave-timeout

```
config igmp snooping leave-timeout <leave_timeout>
```

Description

Configures the IGMP snooping leave timeout.

Syntax Description

leave_timeout	Specifies an IGMP leave timeout value in milliseconds.

Default

1000 ms.

Usage Guidelines

The range is 100 - 100000 ms (10 seconds). If you use sub-second range, you must implement the value in steps of 100 ms. After one second, the value must be implemented in absolute seconds. The specified time is the maximum leave timeout value. After an IGMP leave message is received, the timeout occurs at or before the leave timeout value.

Example

The following command configures the IGMP snooping leave timeout:

```
config igmp snooping leave-timeout 10000
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on “i” series platforms.

config igmp snooping timer

```
config igmp snooping timer <router_timeout> <host_timeout>
```

Description

Configures the IGMP snooping timers.

Syntax Description

router_timeout	Specifies the time in seconds between router discovery.
host_timeout	Specifies the time in seconds between host reports

Default

The router timeout default setting is 260 seconds. The host timeout setting is 260 seconds.

Usage Guidelines

Timers should be set to approximately 2.5 times the router query interval in use on the network. Specify the following:

- router_timeout—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.
- host_timeout—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.
- leave_timeout—A user defined IGMP leave timeout value. The range is 100-100000 ms. The default setting is 1000 ms.

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. An optional optimization for IGMP snooping is the strict recognition of multicast routers only if the remote devices has joined the DVMRP (224.0.0.4) or PIM (244.0.0.13) multicast groups.

Example

The following command configures the IGMP snooping timers:

```
config igmp snooping timer 600 600
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

config pim add vlan

```
config pim add vlan [<name> | all] {dense | sparse}
```

Description

Enables PIM on an IP interface.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.
dense	Specifies PIM dense mode (PIM-DM).
sparse	Specifies PIM sparse mode (PIM-SM).

Default

Dense.

Usage Guidelines

When an IP interface is created, per-interface PIM configuration is disabled by default.

The switch supports both dense mode and sparse mode operation. You can configure dense mode or sparse mode on a per-interface basis. After they are enabled, some interfaces can run dense mode, while others run sparse mode.

Example

The following command enables PIM-DM multicast routing on VLAN *accounting*:

```
config pim add vlan accounting dense
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

config pim cbsr

```
config pim cbsr [vlan <name> {priority <priority> | none]
```

Description

Configures a candidate bootstrap router for PIM sparse-mode operation.

Syntax Description

name	Specifies a VLAN name.
priority	Specifies a priority setting. The range is 0 - 255.
none	Specifies to delete a CBSR.

Default

The default setting is 0 and indicates the lowest priority.

Usage Guidelines

The VLAN specified for CBSR must have ipconfig enabled for PIM sparse mode.

Example

The following command configures a candidate bootstrap router on the VLAN *accounting*:

```
config pim cbsr vlan accounting 30
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all platforms.

config pim crp static

```
config pim crp static <rp address> [none | <access profile>] {<priority>}
```

Description

Configures an RP address and its associated groups statically, for PIM sparse mode operation.

Syntax Description

rp address	Specifies a rendezvous point address.
none	Specifies to check the access profile for permit and deny entries.
access profile	Specifies an access profile name.
priority	Specifies a priority setting. The range is 0 - 255.

Default

The RP is selected dynamically.

Usage Guidelines

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. When 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.

Using PIM-SM, the router sends a join message to the rendezvous point (RP). The RP is a central multicast router that is responsible for receiving and distributing multicast packets. If you use a static RP, all switches in your network must be configured with the same RP address.

Example

The following command configures an RP address and its associated groups statically:

```
config pim crp 10.0.3.1 HQ_10_0_3
```

History

This command was first available in ExtremeWare 6.1.5.

Platform Availability

This command is available on all platforms.

config pim crp timer

```
config pim crp timer <crp_adv_interval>
```

Description

Configures the candidate rendezvous point advertising interval.

Syntax Description

crp_adv_interval	Specifies a candidate rendezvous point advertising interval in seconds.

Default

The default is 60 seconds.

Usage Guidelines

None.

Example

The following command configures the candidate rendezvous point advertising interval to 120 seconds:

```
config pim crp timer 120
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all platforms.

config pim crp vlan access-policy

```
config pim crp vlan <name> access-policy <access_policy> {<priority>}
```

Description

Configures the candidate rendezvous point for PIM sparse-mode operation.

Syntax Description

name	Specifies a VLAN name.
access_policy	Specifies an access policy name.
priority	Specifies a priority setting. The range is 0 - 255.

Default

The default setting is 0 and indicates the highest priority.

Usage Guidelines

The access policy contains the list of multicast group accesses serviced by this RP. To delete a CRP, use the keyword none as the access policy.

The VLAN specified for CBSR must have ipconfig enabled for PIM sparse mode.

Example

The following command configures the candidate rendezvous point for PIM sparse-mode operation on the VLAN *HQ_10_0_3*:

```
config pim crp HQ_10_0_3 rp-list 30
```

History

This command was available in ExtremeWare 6.2.

Platform Availability

This command is available on all platforms.

config pim delete vlan

```
config pim delete vlan [<name> | all]
```

Description

Disables PIM on an interface.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables PIM-DM on VLAN *accounting*:

```
config pim delete vlan accounting
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

config pim register-rate-limit-interval

```
config pim register-rate-limit-interval <time>
```

Description

Configures the initial PIM periodic register rate in case register-stops are not received.

Syntax Description

<time>	Specifies an interval time in seconds. Range is 0 - 60. Default is 0.

Default

Default is 0.

Usage Guidelines

Configuring a non-zero interval time can reduce the CPU load in case register-stops are not received normally.

Example

The following command configures the initial PIM register rate limit interval:

```
config pim register-rate-limit-interval 2
```

History

This command was first available in ExtremeWare 6.2.2.

Platform Availability

This command is available on all platforms.

config pim register-suppress-interval register-probe-interval

```
config pim register-suppress-interval <time> register-probe-interval <time>
```

Description

Configures an interval for periodically sending null-registers.

Syntax Description

register-suppress-interval <time>	Specifies an interval time in seconds. Range is 30 - 200 seconds. Default is 60.
register-probe-interval <time>	Specifies an interval time in seconds. Default is 5.

Default

The following defaults apply:

- register-suppress-interval—60
- register-probe-interval—5

Usage Guidelines

The register-probe-interval time should be set less than the register-suppress-interval time. By default, a null register is sent every 55 seconds (*register-suppress-interval - register-probe-interval*). By specifying an interval, CPU peak load can be avoided because the null-registers are generated in a manageable frequency and time span.

Example

The following command configures the register suppress interval and register probe time:

```
config pim register-suppress-interval 90 register-probe time 10
```

History

This command was first available in ExtremeWare 6.2.2.

Platform Availability

This command is available on all platforms.

config pim register-checksum-to

```
config pim register-checksum-to [include-data | exclude-data]
```

Description

Configures the checksum mode to either include data (for compatibility with Cisco Systems products) or to exclude data (for RFC-compliant operation).

Syntax Description

include-data	Specifies to include data.
exclude-data	Specifies to exclude data.

Default

Include data

Usage Guidelines

None.

Example

The following command configures the checksum mode to include data for compatibility with Cisco Systems products:

```
config pim register-checksum-to include-data
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all platforms.

config pim spt-threshold

```
config pim spt-threshold <leaf_threshold> <rp_threshold>
```

Description

Configures the threshold, in kbps, for switching to SPT. On leaf routers, this setting is based on data packets. On the RP, this setting is based on register packets.

Syntax Description

leaf_threshold	Specifies a leaf threshold.
rp_threshold	Specifies an RP threshold.

Default

The default setting is 0.

Usage Guidelines

For the best performance leveraged by hardware forwarding, use default value "0,0", or small values below 16.

Example

The following command sets the threshold for switching to SPT:

```
config pim spt-threshold 16 8
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all platforms.

config pim timer vlan

```
config pim timer <hello_interval> <jp_interval> vlan [<vlan> | all]
```

Description

Configures the global PIM timers.

Syntax Description

hello_interval	Specifies the amount of time before a hello message is sent out by the PIM router. The range is 1 to 65,519 seconds.
jp_interval	Specifies the join/prune interval. The range is 1 to 65,519 seconds.
vlan	Specifies a VLAN name.
all	Specifies all VLANs.

Default

- hello_interval—30 seconds.
- jp_interval—60 seconds.

Usage Guidelines

None.

Example

The following command configures the global PIM timers on the VLAN *accounting*:

```
config pim timer 150 300 vlan accounting
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

config pim vlan trusted-gateway

```
config pim vlan [<name> | all] trusted-gateway [<access_profile> | none]
```

Description

Configures a trusted neighbor policy.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.
access_profile	Specifies an access profile name.
none	Specifies to check the access profile for permit and deny attributes.

Default

N / A.

Usage Guidelines

Because PIM leverages the unicast routing capability that is already present in the switch, the access policy capabilities are, by nature, different. If the PIM protocol is used for routing IP multicast traffic, the switch can be configured to use an access profile to determine trusted PIM router neighbors for the VLAN on the switch running PIM.

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. When 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.

Example

The following command configures a trusted neighbor policy on the VLAN *backbone*:

```
config pim vlan backbone trusted-gateway nointernet
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

disable dvmrp

```
disable dvmrp
```

Description

Disables DVMRP on the system.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables DVMRP on the system:

```
disable dvmrp
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

disable dvmrp rxmode vlan

```
disable dvmrp rxmode vlan [<name> | all]
```

Description

Disables the receive capability of DVMRP packets on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables the receive capability of DVMRP packets on the VLAN *accounting*:

```
disable dvmrp rxmode vlan accounting
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all platforms.

disable dvmrp txmode vlan

```
disable dvmrp txmode vlan [vlan <name> | all]
```

Description

Disables the transmit capability of DVMRP packets on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

None.

Example

The following command disables the transmit capability of DVMRP packets on the VLAN *accounting*:

```
disable dvmrp txmode vlan accounting
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all platforms.

disable igmp

```
disable igmp {vlan <name>}
```

Description

Disables IGMP on a router interface. If no VLAN is specified, IGMP is enabled on all router interfaces.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

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.

Example

The following command disables IGMP on VLAN *accounting*:

```
disable igmp vlan accounting
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

disable igmp snooping

```
disable igmp snooping {forward-mcrouter-only}
```

Description

Disables IGMP snooping on the switch. If `forward-mcrouter-only` is specified, the switch forwards all multicast traffic to the multicast router only. Otherwise, the switch forwards all multicast traffic to any IP router.

Syntax Description

forward-mcrouter-only	Specifies that the switch forwards all multicast traffic to the multicast router only.

Default

N/A.

Usage Guidelines

Disables IGMP snooping on the whole switch. Two options are available:

- If you do not specify `forward-mcrouter-only`, then the snooping switch will forward all multicast traffic to any IP router (multicast or not).
- If you specify `forward-mcrouter-only`, then the snooping switch will forward all multicast traffic to the multicast router (that is, the router running PIM or DVMRP).

Example

The following command disables IGMP snooping on the switch:

```
disable igmp snooping
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

disable igmp snooping with-proxy

```
disable igmp snooping with-proxy
```

Description

Disables the IGMP snooping proxy. If snooping is not enabled, enabling the proxy also enables snooping. The default setting is enabled.

Enabling the proxy allows the switch to suppress the duplicate join requests on a group to forward to the connected layer 3 switch. The proxy also suppresses unnecessary IGMP leave messages so that they are forwarded only when the first group member joins or the last member leaves the group.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command can be used for troubleshooting purpose. It should be enabled for normal network operation.

Example

The following command disables the IGMP snooping proxy:

```
disable igmp snooping with-proxy
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all platforms.

disable ipmcforwarding

```
disable ipmcforwarding {vlan <name>}
```

Description

Disables IP multicast forwarding on an IP interface.

Syntax Description

name	Specifies a VLAN name.

Default

Disabled.

Usage Guidelines

If no options are specified, all configured IP interfaces are affected. When new IP interfaces are added, IPMC forwarding is disabled by default.

IP forwarding must be enabled before enabling IPMC forwarding, and IPMC forwarding must be disabled before disabling IP forwarding.

Example

The following command disables IPMC forwarding on the VLAN *accounting*:

```
disable ipmcforwarding vlan accounting
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

disable pim

```
disable pim
```

Description

Disables PIM on the system.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command disables PIM on the system:

```
disable pim
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

enable dvmrp

```
enable dvmrp
```

Description

Enables DVMRP on the system.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables DVMRP on the system:

```
enable dvmrp
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

enable dvmrp rxmode vlan

```
enable dvmrp rxmode vlan [<name> | all]
```

Description

Enables the receive capability of DVMRP packets on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

None.

Example

The following command enables the receive capability of DVMRP packets on the VLAN *accounting*:

```
enable dvmrp rxmode vlan accounting
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all platforms.

enable dvmrp txmode vlan

```
enable dvmrp txmode vlan [vlan <name> | all]
```

Description

Enables the transmit capability of DVMRP packets on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

None.

Example

The following command enables the transmit capability of DVMRP packets on the VLAN *accounting*:

```
enable dvmrp txmode vlan accounting
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all platforms.

enable igmp

```
enable igmp {vlan <name>}
```

Description

Enables IGMP on a router interface. If no VLAN is specified, IGMP is enabled on all router interfaces.

Syntax Description

name	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

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.

Example

The following command enables IGMP on the VLAN *accounting*:

```
enable igmp vlan accounting
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

enable igmp snooping

```
enable igmp snooping {forward-mcrouter-only}
```

Description

Enables IGMP snooping on the switch. If `forward-mcrouter-only` is specified, the switch forwards all multicast traffic to the multicast router only. Otherwise, the switch forwards all multicast traffic to any IP router.

Syntax Description

forward-mcrouter-only	Specifies that the switch forwards all multicast traffic to the multicast router only.

Default

N/A.

Usage Guidelines

Enables IGMP snooping on the whole switch. Two options are available:

- If you do not specify `forward-mcrouter-only`, then the snooping switch will forward all multicast traffic to any IP router (multicast or not).
- If you specify `forward-mcrouter-only`, then the snooping switch will forward all multicast traffic to the multicast router (that is, the router running PIM or DVMRP).

Example

The following command enables IGMP snooping on the switch:

```
enable igmp snooping
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

enable igmp snooping with-proxy

```
enable igmp snooping with-proxy
```

Description

Enables the IGMP snooping proxy. If snooping is not enabled, enabling the proxy also enables snooping.The default setting is enabled.

Enabling the proxy allows the switch to suppress the duplicate join requests on a group to forward to the connected layer 3 switch. The proxy also suppresses unnecessary IGMP leave messages so that they are forwarded only when the first group member joins or the last member leaves the group.

Syntax Description

This command has no arguments or variables.

Default

Enabled.

Usage Guidelines

This command can be used for troubleshooting purpose. It should be enabled for normal network operation.

Example

The following command enables the IGMP snooping proxy:

```
enable igmp snooping with-proxy
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all platforms.

enable ipmcforwarding

```
enable ipmcforwarding {vlan <name>}
```

Description

Enables IP multicast forwarding on an IP interface.

Syntax Description

name	Specifies a VLAN name.

Default

Disabled.

Usage Guidelines

If no options are specified, all configured IP interfaces are affected. When new IP interfaces are added, IPMC forwarding is disabled by default.

IP forwarding must be enabled before enabling IPMC forwarding, and IPMC forwarding must be disabled before disabling IP forwarding.

Example

The following command enables IPMC forwarding on the VLAN *accounting*:

```
enable ipmcforwarding vlan accounting
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

enable pim

```
enable pim
```

Description

Enables PIM on the system.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

None.

Example

The following command enables PIM on the system:

```
enable pim
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

run ipmcfdb-check

```
run ipmcfdb-check [index <bucket> <entry> | <ip-multicast-group>
<ip-source-address> vlan <vlan-name>] {extended} {detail}
```

Description

Checks IP multicast FDB entries for consistency.

Syntax Description

bucket	Specifies the bucket portion of the FDB hash index.
entry	Specifies the entry portion of the FDB hash index.
ip-multicast-group	Specifies a multi-cast group. FDB entries with this group will be checked.
ip-source-address	Specifies an IP source address.
vlan-name	Specifies a VLAN name. FDB entries for this VLAN with the specified multicast group number will be checked.
extended	Enables OTP index checking in the MAC entry and VPST of the egress port.
detail	Specifies that more detailed debug information should be logged.

Default

N/A.

Usage Guidelines

The FDB error checking function logs the error count into the system log. Use the `detail` option to log more detailed debug information.

Example

The following command runs a consistency check on the FDB entries for the IP multicast group 168.192.2.4:

```
run ipmcfdb-check 168.192.2.4 195.1.1.100 vlan lab1 extended detail
```

History

This command was first available in ExtremeWare 6.1.9.

Platform Availability

This command is available on all "*i*" series platforms.

show dvmrp

```
show dvmrp {vlan <name> | route {detail}}
```

Description

Displays the DVMRP configuration and statistics, or the unicast route table.

Syntax Description

name	Specifies a VLAN name.
route	Specifies a route.
detail	Specifies to display the information in detailed format.

Default

All.

Usage Guidelines

None.

Example

The following command displays the DVMRP configuration and statistics for the VLAN *accounting*:

```
show dvmrp vlan accounting
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

show igmp group

```
show igmp group [<group_address> [sender [<sender_address>]] [vlan <name>]
```

Description

Lists the IGMP group membership for the specified VLAN.

Syntax Description

group_address	Specifies a group IP address.
sender_address	Specifies a sender's IP address.
name	Specifies a VLAN name.

Default

N / A.

Usage Guidelines

If no VLAN is specified all VLANs are displayed. You can also filter the display by group address and by multicast stream sender address.

Example

The following command lists the IGMP group membership for the VLAN *accounting*:

```
show igmp group 10.0.0.1 10.0.0.2 vlan accounting
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all platforms.

show igmp snooping

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

Description

Displays IGMP snooping registration information and a summary of all IGMP timers and states.

Syntax Description

name	Specifies a VLAN name.
detail	Specifies to display the information in detailed format.

Default

N/A.

Usage Guidelines

The two types of IGMP snooping entry are sender entry and subscribed entry.

The following information is displayed in a per-interface format:

- Group IP multicast address
- Host IP address
- Host VLAN port
- Timeout information

Example

The following command displays IGMP snooping registration information on the VLAN *accounting*:

```
show igmp snooping vlan accounting
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

show ipmc cache

```
show ipmc cache {detail} {<group>} {<source> <netmask>}}
```

Description

Displays the IP multicast forwarding cache.

Syntax Description

detail	Specifies to display the information in detailed format.
group	Specifies an IP group address.
source	Specifies an IP source address.
netmask	Specifies a subnet mask.

Default

N/A.

Usage Guidelines

Displays the following information:

- IP group address
- IP source address / source mask
- Upstream neighbor (RPF neighbor)
- Interface (VLAN-port) to upstream neighbor
- Route expiry time
- Routing protocol
- List of next hop interface (VLAN) and protocol (DVMRP/IGMP)

When the detail option is specified, the switch displays the egress VLAN list and the pruned VLAN list.

Example

The following command displays the IP multicast table for group *224.1.2.3*:

```
show ipmc cache 224.1.2.3
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

show ipmc fdb

```
show ipmc fdb
```

Description

Displays the IP multicast forwarding database.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the IP multicast forwarding database:

```
show ipmc fdb
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all platforms.

show l2stat

```
show l2stat [vlan <name>]
```

Description

Displays the counters for the number of packets bridged, switched, and snooped.

Syntax Description

name	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the counters for the number of packets bridged, switched, and snooped for the VLAN *accounting*:

```
show l2stat accounting
```

History

This command was first available in ExtremeWare 6.2.

Platform Availability

This command is available on all platforms.

show pim

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

Description

Displays the PIM configuration and statistics.

Syntax Description

name	Specifies a VLAN name.

Default

If no VLAN is specified, the configuration is displayed for all PIM interfaces.

Usage Guidelines

None.

Example

The following command displays the PIM configuration and statistics for the VLAN *accounting*:

```
show pim accounting
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

show pim rp-set

```
show pim rp-set {group}
```

Description

Displays the RP-set for one or all groups.

Syntax Description

group	Specifies a group name.

Default

N/A.

Usage Guidelines

None.

Example

The following command displays the RP-set for all groups:

```
show pim rp-set
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all platforms.

unconfig dvmrp

```
unconfig dvmrp {vlan <name>}
```

Description

Resets the DVMRP timers to their default settings.

Syntax Description

name	Specifies a VLAN name.

Default

If no VLAN is specified, all interfaces are reset.

Usage Guidelines

None.

Example

The following command resets all DVMRP timers on VLAN *accounting*:

```
unconfig dvmrp vlan accounting
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

unconfig igmp

```
unconfig igmp
```

Description

Resets all IGMP settings to their default values and clears the IGMP group table.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command resets all IGMP settings to their default values and clears the IGMP group table:

```
unconfig igmp
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

unconfig pim

```
unconfig pim {vlan <name>}
```

Description

Resets all PIM settings on one or all VLANs to their default values.

Syntax Description

name	Specifies a VLAN name.

Default

If no VLAN is specified, the configuration is reset for all PIM interfaces.

Usage Guidelines

None.

Example

The following command resets all PIM settings on the VLAN *accounting*:

```
unconfig pim vlan accounting
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

21 IPX Commands

Basic IPX Command Overview

The switch provides support for the IPX, IPX/RIP, and IPX/SAP protocols. The switch dynamically builds and maintains an IPX routing table and an IPX service table.

The routing software and hardware routes IPX traffic between IPX router interfaces. A router interface is simply a VLAN that has an IPX network identifier (NetID) and IPX encapsulation type assigned to it.

As you create VLANs with different IPX NetIDs the switch automatically routes between them. Both the VLAN switching and IPX routing function occur within the switch.

NOTE

A VLAN can be configured with either an IPX NetID or an IP address. A VLAN cannot be configured for both IPX and IP routing simultaneously.

This chapter describes the IPX commands.

config ipxmaxhops

```
config ipxmaxhops <number>
```

Description

Configures the IPX maximum hop count when forwarding IPX packets.

Syntax Description

number	Specifies a hop count number.

Default

The default setting is 16.

Usage Guidelines

Change the default number only if NetWare Link Services Protocol (NLSP) is running in the IPX network.

Example

The following command configures a maximum IPX hop count of 24:

```
config ipxmaxhops 24
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on the all “*i*” series systems.

config ipxrip add vlan

```
config ipxrip add vlan [<name> | all]
```

Description

Configures one or all IPX VLANs to run IPX/RIP.

Syntax Description

vlan <name>	Specifies a VLAN name.
all	Specifies all VLANs.

Default

IPX/RIP is enabled by default when you configure the IPX VLAN.

Usage Guidelines

IPX/RIP is automatically enabled when a NetID is assigned to the VLAN.

Example

The following command configures IPX VLAN backbone to run IPX/RP:

```
config ipxrip add vlan backbone
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

config ipxrip delete vlan

```
config ipxrip delete vlan [<name> | all]
```

Description

Disables IPX/RIP on one or all interfaces.

Syntax Description

vlan <name>	Specifies a VLAN name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

IPX/RIP is automatically enabled when a NetID is assigned to the VLAN.

Example

The following command disables IPX/RIP on VLAN backbone:

```
config ipxrip delete vlan backbone
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

config ipxrip vlan delay

```
config ipxrip vlan [<name> | all] delay <msec>
```

Description

Configures the time between each IPX/RIP packet within an update interval.

Syntax Description

vlan <name>	Specifies a VLAN name.
all	Specifies all VLANs.
msec	Specifies the delay time in milliseconds.

Default

The default setting is 55 milliseconds.

Usage Guidelines

IPX/RIP is automatically enabled when a NetID is assigned to the VLAN.

Example

The following command configures a delay of 80 milliseconds:

```
config ipxrip vlan accounting delay 80
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

config ipxrip vlan export-filter

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

Description

Assigns an export route filter to an ingress VLAN.

Syntax Description

vlan <name>	Specifies a VLAN name.
all	Specifies all VLANs.
none	Specifies no export filters.
access_profile	Specifies an access profile name.

Default

N / A.

Usage Guidelines

Only the network ID matching the export filter will be added to the IPX route table.

Example

The following command assigns an export route filter to ingress VLAN `accounting`:

```
config ipxrip vlan accounting export-filter my-profile
```

History

This command was introduced in ExtremeWare 4.0; `access-profiles` modified in version 6.1.5b20.

Platform Availability

This command is available on the all "*i*" series systems.

config ipxrip vlan import-filter

```
config ipxrip vlan [<name> | all] import-filter [none | <access_profile>]
```

Description

Assigns an import route filter to an ingress VLAN.

Syntax Description

vlan <name>	Specifies a VLAN name.
all	Specifies all VLANs.
none	Specifies no import filters.
access_profile	Specifies an access profile name.

Default

N / A.

Usage Guidelines

Only the network ID matching the import filter will be added to the IPX route table.

Example

The following command assigns an import route filter to ingress VLAN `accounting`:

```
config ipxrip vlan accounting import-filter my-profile
```

History

This command was introduced in ExtremeWare 4.0; `access-profiles` modified in version 6.1.5b20.

Platform Availability

This command is available on the all "*i*" series systems.

config ipxrip vlan max-packet-size

```
config ipxrip vlan [<name> | all] max-packet-size <size>
```

Description

Configures the maximum transmission unit (MTU) size of the IPX/RIP packet.

Syntax Description

vlan <name>	Specifies a VLAN name.
all	Specifies all VLANs.
size	Specifies the maximum packet size in bytes.

Default

The default setting is 432 bytes.

Usage Guidelines

IPX/RIP is automatically enabled when a NetID is assigned to the VLAN.

Example

The following command configures an MTU size of 128 for the IPX/RIP packet:

```
config ipxrip vlan accounting max-packet-size 128
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

config ipxrip vlan trusted-gateway

```
config ipxrip vlan [<name> | all] trusted-gateway [none | <access_profile>]
```

Description

Assigns an export route filter to the egress VLAN.

Syntax Description

vlan <name>	Specifies a VLAN name.
all	Specifies all VLANs.
none	Specifies no export filters.
access_profile	Specifies an access profile name.

Default

N / A.

Usage Guidelines

Only services matching the trusted gateway are advertised on the egress VLAN.

Example

The following command assigns export route filter `smith` to VLAN `accounting`:

```
config ipxrip vlan accounting trusted-gateway access_profile
```

History

This command was introduced in ExtremeWare 4.0; `access-profiles` modified in version 6.1.5b20.

Platform Availability

This command is available on all platforms.

config ipxrip vlan update-interval

```
config ipxrip vlan [<name> | all] update-interval <time> {hold-multiplier
<number>}
```

Description

Configures the update interval and hold multiplier for IPX/RIP updates.

Syntax Description

vlan <name>	Specifies a VLAN name.
all	Specifies all VLANs.
update-interval <time>	Specifies the update interval time.
hold-multiplier <number>	Specifies the hold multiplier for IPX/RIP updates.

Default

The default update interval is 60 seconds. The default multiplier is 3.

Usage Guidelines

This setting affects both the periodic update interval of IPX/RIP and the aging interval of learned routes. The aging period is calculated using the formula (update-interval * multiplier).

Example

The following command configures the IPX/RIP updates for an update interval of 30 seconds and a hold multiplier of 2 for VLAN accounting:

```
config ipxrip vlan accounting update-interval 30 hold-multiplier 30
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

config ipxroute add

```
config ipxroute add [<dest_netid> | default] <next_hop_id>
<next_hop_node_addr> <hops> <tics>
```

Description

Adds a static IPX route entry in the IPX route table.

Syntax Description

dest_netid	Specifies the destination NetID.
next_hop_id	Specifies the NetID of the neighbor IPX network.
next_hop_node_addr	Specifies the node address of the next IPX router.
hops	Specifies the maximum hop count.
tics	Specifies the timer delay value.

Default

N/A.

Usage Guidelines

Static routes are used to reach networks not advertised by routers. You can configure up to 64 static IPX routes on the switch. Static routes are never aged out of the routing table. Static routes are advertised to the network using IPX/RIP.

Example

The following command adds a static IPX route entry to the IPX route table:

```
config ipxroute add default 0011 00:eb:2a:0b:1e:0a
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

config ipxroute delete

```
config ipxroute delete [<dest_netid> | default] <next_hop_netid>
<next_hop_node_addr>
```

Description

Removes a static IPX route entry from the route table.

Syntax Description

dest_netid	Specifies the destination NetID.
next_hop_id	Specifies the NetID of the neighbor IPX network.
next_hop_node_addr	Specifies the node address of the next IPX router.

Default

N / A.

Usage Guidelines

If you have defined default or static routes, those routes will remain in the configuration independent of whether the VLAN or VLAN IP address that used them remains. You should manually delete the routes if no VLAN IP address is capable of using them.

Example

The following command deletes a static IPX route entry to the IPX route table:

```
config ipxroute delete default 0011 00:eb:2a:0b:1e:0a
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

config ipxsap add vlan

```
config ipxsap add vlan [<name> | all]
```

Description

Configures an IPX VLAN to run IPX/SAP routing.

Syntax Description

vlan <name>	Specifies a VLAN name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

If no VLAN is specified, all VLANs are configured to run IPX/SAP routing. IPX/SAP routing is enabled by default when the IPX VLAN is configured.

Example

The following command configures the IPX VLAN accounting to run IPX/SAP routing:

```
config ipxsap add vlan accounting
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

config ipxsap delete vlan

```
config ipxsap delete vlan [<name> | all]
```

Description

Disables IPX/SAP on an interface.

Syntax Description

vlan <name>	Specifies a VLAN name.
all	Specifies all VLANs.

Default

N/A.

Usage Guidelines

If no VLAN is specified, all VLANs are configured to run IPX/SAP routing. IPX/SAP routing is enabled by default when the IPX VLAN is configured.

Example

The following command disables IPX/SAP on VLAN accounting:

```
config ipxsap delete vlan accounting
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

config ipxsap vlan delay

```
config ipxsap vlan [<name> | all] delay <msec>
```

Description

Configures the time between each SAP packet within an update interval.

Syntax Description

vlan <name>	Specifies a VLAN name.
all	Specifies all VLANs.
msec	Specifies a delay in milliseconds.

Default

The default setting is 55 milliseconds.

Usage Guidelines

If no VLAN is specified, all VLANs are configured to run IPX/SAP routing. IPX/SAP routing is enabled by default when the IPX VLAN is configured.

Example

The following command sets the time between each SAP packet to 40 milliseconds for VLAN accounting:

```
config ipxsap vlan accounting delay 40
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

config ipxsap vlan export-filter

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

Description

Assigns an export route filter to an ingress VLAN.

Syntax Description

vlan <name>	Specifies a VLAN name.
all	Specifies all VLANs.
none	Specifies no export filters
access_profile	Specifies an access profile name.

Default

N / A.

Usage Guidelines

Only the network ID matching the export filter will be added to the IPX route table.

Example

The following command assigns an export route filter to ingress VLAN `accounting`:

```
config ipxsap vlan accounting export-filter none
```

History

This command was introduced in ExtremeWare 4.0; `access-profiles` modified in version 6.1.5b20.

Platform Availability

This command is available on the all "*i*" series systems.

config ipxsap vlan import-filter

```
config ipxsap vlan [<name> | all] import-filter [none | access_profile]
```

Description

Assigns an import route filter to an ingress VLAN.

Syntax Description

vlan <name>	Specifies a VLAN name.
all	Specifies all VLANs.
none	Specifies no route filters.
access_profile	Specifies an access profile name.

Default

N / A.

Usage Guidelines

Only the network ID matching the import filter will be added to the IPX route table.

Example

The following command assigns an import route filter to ingress VLAN accounting:

```
config ipxsap vlan accounting import-filter none
```

History

This command was introduced in ExtremeWare 4.0; access-profiles modified in version 6.1.5b20.

Platform Availability

This command is available on the all "*i*" series systems.

config ipxsap vlan max-packet-size

```
config ipxsap vlan [<name> | all] max-packet-size <number>
```

Description

Configures the MTU size of the IPX/SAP packets.

Syntax Description

vlan <name>	Specifies a VLAN name.
all	Specifies all VLANs.
max-packet-size <number>	Specifies the maximum packet size in bytes.

Default

The default setting is 432 bytes.

Usage Guidelines

If no VLAN is specified, all VLANs are configured to run IPX/SAP routing. IPX/SAP routing is enabled by default when the IPX VLAN is configured.

Example

The following command configures an MTU size of 356 bytes for the IPX/SAP packets on VLAN accounting:

```
config ipxsap vlan [<name> | all] max-packet-size <number>
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

config ipxsap vlan trusted-gateway

```
config ipxsap vlan [<name> | all] trusted-gateway [none | <access_profile>]
```

Description

Assigns an export SAP service filter to the egress VLAN.

Syntax Description

vlan <name>	Specifies a VLAN name.
all	Specifies all VLANs.
none	Specifies no service filters.
access_profile	Specifies an access profile name.

Default

N / A.

Usage Guidelines

Only the services matching the trusted-gateway are advertised on the egress VLAN.

Example

The following command assigns an export SAP service filter named `smith` to VLAN `accounting`:

```
config ipxsap vlan accounting trusted-gateway smith
```

History

This command was introduced in ExtremeWare 4.0; `access-profiles` modified in version 6.1.5b20.

Platform Availability

This command is available on all platforms.

config ipxsap vlan update-interval

```
config ipxsap vlan [<name> | all] update-interval <time> {hold-multiplier
<number>}
```

Description

Configures the update interval and hold multiplier for IPX/SAP updates.

Syntax Description

vlan <name>	Specifies a VLAN name.
all	Specifies all VLANs.
update-interval <time>	Specifies the update interval time.
hold-multiplier <number>	Specifies the hold multiplier for IPX/RIP updates.

Default

The default update interval is 60 seconds. The default multiplier is 3.

Usage Guidelines

This setting affects both the periodic update interval of SAP and the aging interval of learned routes. The default update interval is 60 seconds. The aging period is calculated using the formula (update-interval * multiplier). The default multiplier is 3. Triggered update is always enabled; therefore, new information is processed and propagated immediately.

Example

The following command configures an update interval of `30` seconds and a hold multiplier of 2 for the IPX/SAP updates for VLAN `accounting`:

```
config ipxsap vlan accounting update-interval 30 hold-multiplier 2
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

config ipxsap vlan gns-delay

```
config ipxsap vlan <name> gns-delay <msec>
```

Description

Configures the amount of time the switch waits before answering a GNS request.

Syntax Description

vlan <name>	Specifies a VLAN name.
msec	Specifies a delay in milliseconds.

Default

The switch answers a GNS request as soon as possible (0 milliseconds).

Usage Guidelines

If no VLAN is specified, all VLANs are configured to run IPX/SAP routing. IPX/SAP routing is enabled by default when the IPX VLAN is configured.

Example

The following command sets a GNS delay time of `20` milliseconds on VLAN `accounting`:

```
config ipxsap vlan accounting gns-delay 20
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

config ipxservice add

```
config ipxservice add <service_type> <service_name> <netid> <mac_address>
<socket> <hops>
```

Description

Adds a static entry to the IPX service table.

Syntax Description

service_type	Specifies a service type.
service_name	Specifies a service name.
netid	Specifies the IPX network identifier of the server.
mac_address	Specifies the MAC address of the server.
socket	Specifies the IPX port number on the server.
hops	Specifies the number of hops (for SAP routing purposes).

Default

N/A.

Usage Guidelines

Service information may also be entered into the IPX Service Table dynamically, by way of SAP.

The `socket` provides you with access to a particular function on the server.

Example

The following command adds non-advertising server `chalk` to the IPX service table, with `0004` as SAP for a file server, `00:AO:C9:17:22:F5` as the MAC address, `0451` as the socket number for a connection request, and `3` as the number of hops to the server:

```
config ipxservice add chalk 0004 00:AO:C9:17:22:F5 0451 3
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

config ipxservice delete

```
config ipxservice delete <service_type> <service_name> <netid>
<mac_address> <socket>
```

Description

Deletes an IPX service from the service table.

Syntax Description

service_type	Specifies a service type.
service_name	Specifies a service name.
netid	Specifies the IPX network identifier of the server.
mac_address	Specifies the MAC address of the server.
socket	Specifies the IPX port number on the server.

Default

N / A.

Usage Guidelines

The service information is entered into the IPX Service Table in one of the following two ways:

- Dynamically, by way of SAP
- Statically, using the `config ipxservice add` command

The `socket` provides you with access to a particular function on the server.

Example

The following command deletes non-advertising server `chalk` from the IPX service table, with `0004` as SAP for a file server, `00:AO:C9:17:22:F5` as the MAC address, `0451` as the socket number for a connection request, and `3` as the number of hops to the server.

```
config ipxservice delete chalk 0004 00:AO:C9:17:22:F5 0451
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

config vlan xnetid

```
config vlan <name> xnetid <netid> [enet_ii | enet_8023 | enet_8022 |
enet_snap]
```

Description

Configures a VLAN to use a particular encapsulation type.

Syntax Description

vlan <name>	Specifies a VLAN name.
netid	Specifies the IPX network identifier of the server.
enet_ii	Specifies an Ethernet 2 header.
enet_8023	Specifies the IEEE 802.3 length field.
enet_8022	Specifies and IEEE format and includes the IEEE 802.2 LLC header.
enet_snap	Specifies to add SNAP header to the IEEE 802.2 LLC header.

Default

N/A.

Usage Guidelines

Novell NetWare supports four types of frame encapsulation. The ExtremeWare term for each type is shown in the following list:

Table 22:

ENET_II	The frame uses the standard Ethernet 2 header.
ENET_8023	The frame includes the IEEE 802.3 length field, but does not include the IEEE 802.2 Logical Link Control (LLC) header. This encapsulation is used by NetWare version 2.x and the original NetWare 3.x version.
ENET_8022	The frame uses the standard IEEE format and includes the IEEE 802.2 LLC header. This encapsulation is used by NetWare version 3.12 and 4.x.
ENET_SNAP	The frame adds a Subnetwork Access Protocol (SNAP) header to the IEEE 802.2 LLC header.

Example

The following command configures VLAN `Support` to use encapsulation `enet_8022`:

```
config vlan Support xnetid A2B5 enet_8022
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

disable ipxrip

```
disable ipxrip
```

Description

Disables IPX/RIP on the router.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

IPX/RIP is automatically enabled when a NetID is assigned to the VLAN.

Example

The following disables IPX/RIP on the router:

```
disable ipxrip
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

disable ipxsap

```
disable ipxsap
```

Description

Disables IPX/SAP on the router.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

If no VLAN is specified, all VLANs are configured to run IPX/SAP routing. IPX/SAP routing is enabled by default when the IPX VLAN is configured.

Example

The following disables IPX/SAP on the router:

```
disable ipxsap
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

disable ipxsap gns-reply

```
disable ipxsap gns-reply {vlan <name>}
```

Description

Disables Get Nearest Server (GNS) reply on one or all IPX interfaces.

Syntax Description

vlan <name>	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

ExtremeWare supports the GNS reply function. When a NetID is assigned to the switch, the GNS reply service is automatically enabled. When a station requests a particular service on the network (for example, locating a print server), the station sends a GNS request and the switch responds to the request. If GNS-reply is disabled, the switch drops the request.

Example

The following command disables GNS reply on IPX VLAN accounting:

```
disable ipxsap gns-reply vlan accounting
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

disable type20 forwarding

```
disable type20 forwarding {vlan <name>}
```

Description

Disables the forwarding of IPX type 20 packets.

Syntax Description

vlan <name>	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

Type 20 packets are NetBIOS inside IPX.

Example

The following command disables the forwarding of IPX type 20 packets for VLAN `accounting`:

```
disable type20 forwarding vlan accounting
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

enable ipxrip

```
enable ipxrip
```

Description

Enables IPX/RIP on the router.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

IPX/RIP is automatically enabled when a NetID is assigned to the VLAN.

Example

The following command enables IPX/RIP on the router:

```
enable ipxrip
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

enable ipxsap

```
enable ipxsap
```

Description

Enables IPX/SAP on the router.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

If no VLAN is specified, all VLANs are configured to run IPX/SAP routing. IPX/SAP routing is enabled by default when the IPX VLAN is configured.

Example

The following command enables IPX/SAP on the router:

```
enable ipxsap
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

enable ipxsap gns-reply

```
enable ipxsap gns-reply {vlan <name>}
```

Description

Enables GNS reply on one or all IPX interfaces.

Syntax Description

vlan <name>	Specifies a VLAN name.

Default

The default setting is enabled.

Usage Guidelines

If no VLAN is specified, GNS reply is enabled on all IPX interfaces.

Example

The following command enables GNS reply for IPX VLAN accounting:

```
enable ipxsap gns-reply vlan accounting
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

enable type20 forwarding

```
enable type20 forwarding {vlan <name>}
```

Description

Enables the forwarding of IPX type 20 packets.

Syntax Description

vlan <name>	Specifies a VLAN name.

Default

Enabled.

Usage Guidelines

Type 20 packets are NetBIOS inside IPX.

Example

The following command enables the forwarding of IPX type 20 packets for VLAN accounting:

```
enable type20 forwarding vlan accounting
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all “*i*” series systems.

show ipxconfig

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

Description

Displays IPX configuration information for one or all VLANs.

Syntax Description

vlan <name>	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

This command is analogous to the `show ipconfig` command for the IP protocol. It displays summary global IPX configuration information followed by per-VLAN information. Information includes enable/disable status for IPX/RIP, IPX/SAP, IPX route sharing, IPX service sharing, and so on.

Example

The following command displays the IPX configuration information for VLAN `accounting`:

```
show ipxconfig vlan accounting
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

show ipxfdb

```
show ipxfdb {vlan <name> | xnetid <netid>}
```

Description

Displays the hardware IPX FDB information.

Syntax Description

vlan <name>	Specifies a VLAN name.
netid	Specifies an IPX network number.

Default

N / A.

Usage Guidelines

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

Example

The following command displays the hardware IPX FDB information for VLAN `accounting`:

```
show ipxfdb vlan accounting
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all “*i*” series systems.

show ipxrip

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

Description

Displays IPX/RIP configuration and statistics for one or all VLANs.

Syntax Description

vlan <name>	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

The enable status of IPX/RIP displayed includes operational and administrative status. It also lists any identified IPX/RIP neighbors, RIP packet statistics, and several other timer settings.

Example

The following command displays the IPX/RIP configuration information and statistics for VLAN `accounting`:

```
show ipxrip vlan accounting
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

show ipxroute

```
show ipxroute {vlan <name> | xnetid <netid> | origin [static | rip |
local]}
```

Description

Displays the IPX routes in the route table.

Syntax Description

vlan <name>	Specifies a VLAN name.
netid	Specifies an IPX network number.
static	Specifies a statically defined route.
rip	Specifies a RIP learned route.
local	Specifies a local interface.

Default

N/A.

Usage Guidelines

Route information is entered into the IPX route table in one of the following two ways:

- Dynamically, by way of RIP
- Statically, using the `config ipxroute add` command

IPX/RIP is automatically enabled when a NetID is assigned to the VLAN. To remove the advertisement of an IPX VLAN, use the `config ipxrip delete` command.

Example

The following command displays the IPX routes in the route table for VLAN `accounting`:

```
show ipxroute vlan accounting
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

show ipxsap

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

Description

Displays the enable status of IPX/SAP for the VLAN, and its operational and administrative status (including the GNS reply service). It also lists any identified IPX/SAP neighbors, SAP packet statistics, and several other timer settings.

Syntax Description

vlan <name>	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

None.

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

show ipxservice

```
show ipxservice {vlan <name> | xnetid <netid> _ origin [static | sap |
local]}
```

Description

Displays IPX services learned by way of SAP.

Syntax Description

vlan <name>	Specifies a VLAN name.
netid	Specifies an IPX network number.

Default

N / A.

Usage Guidelines

The service information is entered into the IPX Service Table in one of the following two ways:

- Dynamically, by way of SAP
- Statically, using the `config ipxservice add` command

Example

The following command displays IPX/SAP service information for VLAN `accounting`:

```
show ipxservice vlan accounting
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

show ipxstats

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

Description

Displays IPX packet statistics for the IPX router, and one or all VLANs.

Syntax Description

vlan <name>	Specifies a VLAN name.

Default

All VLANs.

Usage Guidelines

Displays both RIP and SAP packet statistics.

Example

The following command displays IPX packet statistics for VLAN accounting:

```
show ipxstats vlan accounting
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

unconfig ipxrip

```
unconfig ipxrip {vlan <name>}
```

Description

Resets the IPX/RIP settings on one or all VLANs to the default.

Syntax Description

vlan <name>	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

Removes import and export filters, and resets the MTU size, update interval, and inter-packet delay.

Example

The following command

```
unconfig ipxrip vlan backbone
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

unconfig ipxsap

```
unconfig ipxsap {vlan <name>}
```

Description

Resets the IPX/SAP settings on one or all VLANs to the default.

Syntax Description

vlan <name>	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

Removes import and export filters, and resets the MTU size, update interval, and inter-packet delay.

Example

The following command resets the IPX/SAP settings on VLAN backbone to the defaults:

```
unconfig ipxsap vlan backbone
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

unconfig vlan xnetid

```
unconfig vlan <name> xnetid
```

Description

Removes the IPX NetID of a VLAN.

Syntax Description

vlan <name>	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

IPX/RIP is automatically enabled when a NetID is assigned to the VLAN.

Example

The following command removes the IPX NetID of VLAN accounting:

```
unconfig vlan accounting xnetid
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

xping

```
xping {continuous} {size <n>} <netid> <node_address>
```

Description

Pings an IPX node specified by the network ID and the node address.

Syntax Description

continuous	Specifies that pings are to be sent continuously.
size <n>	Specifies the ping packet size in bytes.
netid	Specifies an IPX network number.

Default

N / A.

Usage Guidelines

If `continuous` is not specified, four pings are sent. The default ping packet size is 256 data bytes. The size range is between 1 and 1,484 bytes.

Example

The following command pings IPX node `0010460` with a node address of `00:2b:2a:00:1c:0a`:

```
xping 0010460 00:2b:2a:00:1c:0a
```

History

This command was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on the all "*i*" series systems.

22 ARM Commands

The Accounting and Routing Module (ARM) is a self-contained module for the BlackDiamond 6800 series chassis-based system. Unlike other BlackDiamond modules, there are no external network interfaces on the ARM. Instead, the ARM provides advanced IP services for the other input/output (I/O) modules installed in the chassis. The ARM contains a powerful set of packet processing resources that operate in a one-armed fashion: receiving frames from the switch fabric, processing the frames, and transmitting the frames back into the switch fabric. More specifically, the accounting feature is used to track and record IP unicast packets. This enables you to create custom billing rates for your customers.

The two main features of the ARM are: IP unicast forwarding and destination-sensitive accounting.

In the first feature, IP unicast packets are routed in the hardware using the longest prefix match algorithm. Counts of packets and bytes are maintained based on the routes used to forward packets.

In the second feature, destination-sensitive accounting collects statistics that are maintained for forwarded IP traffic to support billing on a destination basis. To configure destination-sensitive accounting, a bin number can be assigned to one or more IP route entries using the ExtremeWare `route-map` command.

Bin numbers are integers that range from 0-7 and their only intrinsic meaning is to identify a particular set of accounting statistics. Each bin contains a 64-bit count of the number of packets that have been forwarded and a 64-bit count of the number of bytes that have been forwarded. When the MPLS module forwards an IP packet, the bin number from the forwarding database entry for the IP destination is used to identify the set of counters to be updated.

Eight unique bins are maintained for each of the possible 4096 VLAN IDs. Logically, the bins are organized as a two-dimensional array, with the row index being a VLAN ID and the column index being a bin number. Thus, when an IP frame is forwarded, the input VLAN ID selects the row and the bin number from the forwarding database entry selects the column. The use of input VLAN ID enables billing statistics to be maintained on a per customer basis where the VLAN ID identifies the customer.

This chapter documents the ARM command set. Some commands are new for the ARM; other commands have been enhanced to support the ARM.

Basic Accounting Configuration Information

This section uses several typical usage and configuration schemes to provide a brief overview of the destination-sensitive accounting configuration process as a general context for the detailed command description sections that follow.

In the most basic terms, to enable the accounting function, you must enable the accounting feature, create a customer VLAN ID, enable IP forwarding, and configure the accounting bin using the route map feature.

You use a special set of commands to configure the MPLS module to initiate the accounting function.

NOTE

Support for MPLS modules is included in an ExtremeWare IP Services Technology Release, currently based on ExtremeWare v6.1.8b12. Later versions of ExtremeWare (6.1.9 or 6.2) currently do not support MPLS modules.

clear accounting counters

```
clear accounting counters
```

Description

Clears (zeroes out) all of the billing statistics.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command clears (zeroes out) all of the billing statistics.:

```
clear accounting counters
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config route-map set accounting-index 1 value

```
config route-map <route-map> <sequence_number> [add | delete] set
accounting-index 1 value <bin_number>
```

Description

Configures the accounting bin number to be associated with the specified route map entry.

Syntax Description

route-map	Specifies a route map.
sequence number	Specifies a specific entry in the route map.
add	Specifies to add the statement to the route map.
delete	Specifies to delete the statement from the route map.
bin_number	Specifies an accounting bin number.

Default

N / A.

Usage Guidelines

- The accounting-index value is always set to 1 for destination-sensitive accounting.
- The `route-map` parameter identifies a particular route map.
- The `sequence_number` parameter identifies a specific entry in that route map. The sequence number must be associated with a match statement.
- The `set accounting-index 1 value` keyword phrase indicates that the following parameter is an accounting bin number.
- The `bin_number` parameter is an integer between 0—7, and allows you to define the accounting bin number.

Table 23: Set Operation Keywords

Command	Description of Change
accounting-index <index> value <value>	Sets the accounting bin number for the route-mapped accounting index. The accounting index value is always set to 1 for destination-sensitive accounting.

Example

The following command configures the accounting bin number to be associated with the specified route map entry:

```
config route-map rt40 11 add set accounting-index 1 value 5
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

disable accounting

```
disable accounting
```

Description

Disables the destination-sensitive accounting function.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Destination-sensitive accounting collects statistics that are maintained for forwarded IP traffic to support billing on a destination basis. To configure destination-sensitive accounting, a bin number can be assigned to one or more IP route entries using the ExtremeWare `route-map` command.

Example

The following command disables the destination-sensitive accounting function:

```
disable accounting
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

enable accounting

```
enable accounting
```

Description

Enables the destination-sensitive accounting function.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Destination-sensitive accounting collects statistics that are maintained for forwarded IP traffic to support billing on a destination basis. To configure destination-sensitive accounting, a bin number can be assigned to one or more IP route entries using the ExtremeWare `route-map` command.

Example

The following command enables the destination-sensitive accounting function:

```
enable accounting
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

show accounting

```
show accounting {<vlan>}
```

Description

Displays accounting statistics for the specified VLAN. If no VLAN is specified, statistics for all VLANs are displayed.

Syntax Description

vlan	Specifies a VLAN name.

Default

N/A.

Usage Guidelines

You can display the accounting statistics for a single VLAN or all VLANs by issuing the `show accounting <vlan>` command. The `show accounting <vlan>` command lists the packet and octet counts for each bin number per VLAN. Omitting the VLAN name displays the accounting statistics for all the VLANs.

Example

The following command displays accounting statistics for the *vlan1* VLAN:

```
show accounting vlan1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

23 MPLS Commands

The MultiProtocol Label Switching (MPLS) module is a self-contained module for the BlackDiamond 6800 series chassis-based system. Unlike other BlackDiamond modules, there are no external network interfaces on the MPLS module. Instead, the MPLS module provides advanced IP services for the other input/output (I/O) modules installed in the chassis. The MPLS module contains a powerful set of packet processing resources that operate in a one-armed fashion: receiving frames from the switch fabric, processing the frames, and transmitting the frames back into the switch fabric.

MPLS encompasses a growing set of protocols defined by the IETF. True to its name, MPLS is based on a label-switching forwarding algorithm. ATM and Frame Relay are examples of other protocols that use label-switching forwarding algorithms.

Conceptually, label switching is straightforward. A label is a relatively short, fixed-length identifier that is used to forward packets received from a given link. The label value is locally significant to a particular link and is assigned by the receiving entity.

Because labels are relatively short (for example, 20 bits in a MPLS shim header), the label of a received packet can be used as an index into a linear array containing the forwarding database. Forwarding database entries indicate the outgoing port and any label(s) to be applied to forwarded frames. Thus, forwarding may consist of a simple lookup and replacement of the incoming label with the appropriate outgoing label (otherwise known as *label swapping*).

This chapter documents the MPLS command set. Some commands are new for the MPLS module; other commands have been enhanced to support the MPLS module.

NOTE

Support for MPLS modules is included in an ExtremeWare IP Services Technology Release, currently based on ExtremeWare v6.1.8b12. Later versions of ExtremeWare (6.1.9 or 6.2) currently do not support MPLS modules.

config mpls

```
config mpls [ldp | targeted-ldp] [hello | keep-alive] <hold_time>
<interval_time>
```

Description

Configures LDP session timers.

Syntax Description

ldp	Specifies an LDP session.
targeted-ldp	Specifies a targeted LDP session.
hello <hold_time> <interval_time>	The amount of time (in seconds) that a hello message received from a neighboring LSR remains valid. If a hello message is not received from a particular neighboring LSR within the specified `hello <hold_time>`, the hello-adjacency is not maintained with that neighboring LSR. The range is 6 to 65,534.
keep-alive <hold_time> <interval_time>	The time (in seconds) during which an LDP message must be received for the LDP session with a particular peer LSR to be maintained. If an LDP PDU is not received within the specified session `keep-alive <interval_time>`, the corresponding LDP session is torn down.The `<hold_time>` range is 6 to 65,534. The `<interval_time>` range is 1 to 21844.

Default

ldp hello <hold_time> – 15 seconds

targeted-ldp hello <hold_time> – 45 seconds

ldp hello <interval_time> – 5 seconds

targeted-ldp hello <interval_time> – 15 seconds

ldp keep-alive <hold_time> – 40 seconds

targeted-ldp keep-alive <hold_time> – 60 seconds

ldp keep-alive <interval_time> – 13 seconds

targeted-ldp keep-alive <interval_time> – 20 seconds

Usage Guidelines

LDP session timers are separately configurable for LDP and targeted LDP sessions. The `hello <hold_time> <interval_time>` parameter specifies the amount of time (in seconds) that a hello message received from a neighboring LSR remains valid. If a hello message is not received from a particular neighboring LSR within the specified `hello <hold_time>`, the hello-adjacency is not maintained with that neighboring LSR.

The session `keep-alive <hold_time> <interval_time>` parameter specifies the time (in seconds) during which an LDP message must be received for the LDP session with a particular peer LSR to be

maintained. If an LDP PDU is not received within the specified session `keep-alive <interval_time>`, the corresponding LDP session is torn down.

The minimum and maximum values for both the `hello <hold_time> <interval_time>` and `keep-alive <hold_time> <interval_time>` are 6 and 65,534, respectively.

This command can only be executed when MPLS is disabled.

Example

The following command configures LDP session hello hold time to 30 seconds and the interval time to 5 seconds:

```
config mpls ldp hello 30 5
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config mpls add tls-tunnel

```
config mpls add tls-tunnel <tunnel_name> [lsp <lsp_name> | <ipaddress> |
<host_name>] <local_vlan_name> [tls-labels <ingress_label> <egress_label> |
vcid <vcid> {<groupid>} {from [<local_endpoint_ipaddress> |
<local_endpoint_vlan>]}]
```

Description

Adds a TLS tunnel.

Syntax Description

tunnel_name	Specifies a name used to identify the TLS tunnel within the switch.
[lsp <lsp_name> | <ipaddress> | <host_name>]	Identifies the peer LSR that is the tunnel endpoint. The DNS client must be configured to use the `<host_name>`.
local_vlan_name	Specifies a VLAN name that identifies the layer 2 traffic that is to be transported.
tls-labels <ingress_label> <egress_label>	Identifies the innermost labels of the tunnel stack.
vcid	Identifies the virtual circuit identifier. The `vcid` value is a non-zero, 32-bit number.
groupid	Identifies the logical VCID group number. The `groupid` is a 32-bit number. All TLS tunnels that are members of the same TLS group ID can be withdrawn simultaneously by specifying the `groupid`.
from <local_endpoint_ipaddress> | <local_endpoint_vlan>	Identifies the local endpoint of the TLS tunnel.

Default

N/A.

Usage Guidelines

To add a static labeled TLS tunnel, use the following command:

```
config mpls add tls-tunnel <tunnel_name> [lsp <lsp_name> | <ipaddress> | <host_name>]
<local_vlan_name> tls-labels <ingress_label> <egress_label>
```

To add a dynamic labeled TLS tunnel (martini-draft compliant), use the following command:

```
config mpls add tls-tunnel <tunnel_name> [lsp <lsp_name> | <ipaddress> | <host_name>]
<local_vlan_name> vcid <vcid> <groupid>
```

The `<tunnel_name>` parameter is a character string that is to be used to identify the TLS tunnel within the switch. It must begin with an alphabetic character and can contain up to 31 additional alphanumeric characters.

The `<ipaddress>` parameter identifies the peer LSR that is the endpoint of the tunnel. This IP address should be configured with a 32-bit prefix on the peer LSR. When the peer LSR is also an Extreme switch, either OSPF must also be enabled on the VLAN to which the IP address is assigned (using the

config ospf add vlan command on the peer switch), or the peer switch must be configured to distribute direct routes into the OSPF domain (using the enable ospf export direct command). The ospf export command should be used when the tunnel LSP needs to cross OSPF area boundaries or when ESRP is enabled on the VLAN to which the IP address is assigned.

The <vcid> parameters are used to configure dynamic TLS tunnels when full martini-draft TLS tunnel compliance is desired. The vcid and groupid values are advertised on a targeted LDP session to the specified tunnel endpoint ipaddress in a martini-draft defined FEC-TLV. Each LER advertises the vcid, groupid, and VLAN label in the Label Mapping message across an LDP session. This three-tuple TLS tunnel information allows each egress LER to dynamically bind the TLS tunnel to a local VLAN. The vcid is a non-zero 32-bit ID that defines the tunnel connection and the optionally specified groupid is a 32-bit value that defines logical virtual tunnel connection group. The groupid value defaults to zero if not explicitly configured.

The <local_vlan_name> parameter identifies the Layer-2 traffic that is to be transported. All of the local traffic received by the switch for this VLAN is transported across the tunnel.

The tls-labels parameters specify the innermost labels of the tunnel label stack and are used to configure static TLS label tunnels. The <egress_label> is inserted into the MPLS header of Layer-2 frames forwarded onto the tunnel LSP by this switch, and must be meaningful to the peer TLS node.

All traffic received from the tunnel LSP that contains the <ingress_label> is forwarded to the local VLAN identified by the <local_vlan_name> parameter.

When ingress traffic is forwarded to the local VLAN, the VLAN ID is set to the VLAN ID of the local VLAN, without regard to the VLAN ID in the MAC header of the frame received from the tunnel LSP. Thus, there is no requirement that all sites of an extended VLAN be configured to use the same VLAN ID. This can simplify network management in some situations.

The tls-labels parameters are specified using hexadecimal notation. The value of the <ingress_label> parameter must be unique within the switch (the same <ingress_label> value cannot be used for two different tunnels). The valid range of the ingress label parameter is [8C000..8FFFF].

The valid range of the <egress_label> parameter is [00010..FFFFF]. If the peer LSR is also an Extreme switch, then the <egress_label> must be in the range [8C000..8FFFF].

Because LSPs are unidirectional in nature, coordinated configuration is required at both tunnel endpoint switches. The <egress_label> at one tunnel endpoint switch must match the <ingress_label> at the other tunnel endpoint switch, and vice versa.

Example

The following command creates a TLS tunnel to 11.0.4.11 for traffic originating from VLAN unc:

```
config mpls add tls-tunnel rt40 11.0.4.11 unc tls-labels 8f001 8f004
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config mpls add vlan

```
config mpls add vlan [<name> | all] {ldp | rsvp-te}
```

Description

Enables LDP or RSVP-TE for one or all VLANs.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.
ldp	Enables LDP.
rsvp-te	Enables RSVP-TE.

Default

N / A.

Usage Guidelines

MPLS must be enabled on all VLANs that transmit or receive MPLS-encapsulated frames. Using the `config mpls add vlan` command causes the LDP neighbor discovery process to begin on the specified VLAN.

NOTE

The specified VLAN must be configured with an IP address and must have IP forwarding enabled. IGMP snooping must also be enabled on the switch.

If all VLANs are selected, MPLS is enabled on all VLANs that have an IP address and IP forwarding enabled.

If you have enabled MPLS on an OSPF interface that is used to reach a particular destination, make sure that you enable MPLS on all additional OSPF interfaces that can reach that same destination (for example, enable MPLS on all VLANs that are connected to the backbone network).

Example

The following command enables RSVP-TE on vlan1:

```
config mpls add vlan vlan1 rsvp-te
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config mpls delete tls-tunnel

```
config mpls delete tls-tunnel [<tunnel_name> | group <groupid> | all]
```

Description

Deletes one or all TLS tunnels.

Syntax Description

tunnel_name	Specifies a TLS tunnel name.
group <groupid>	Specifies a group identifier
all	Specifies all TLS tunnels.

Default

N / A.

Usage Guidelines

This command deletes the TLS tunnel with the specified tunnel name. Specify the `<groupid>` if you want to delete all TLS tunnels belonging to a specific group. Specify the `<groupid>` if you want to delete all TLS tunnels belonging to a specific group. Use the `all` keyword to delete all TLS tunnels.

Example

The following command deletes the TLS tunnel *rt40*:

```
config mpls delete tls-tunnel rt40
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config mpls delete vlan

```
config mpls delete vlan [<name> | all] {ldp | rsvp-te}
```

Description

Disables LDP or RSVP-TE on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.
ldp	Disables LDP.
rsvp-te	Disables RSVP-TE.

Default

N / A.

Usage Guidelines

Disables LDP or RSVP-TE on one or all VLANs. If not specified, both are disabled for the specified VLAN.

Example

The following command disables RSVP-TE on vlan1:

```
config mpls delete vlan vlan1 rsvp-te
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config mpls ldp advertise

```
config mpls ldp advertise [direct | rip | static] [all | none | route-map
<route_map>]
```

Description

Configures a filter to be used by LDP when originating unsolicited label mapping advertisements to LDP neighbors.

Syntax Description

direct	Specifies that the advertisement filter is applied to the associated FECs with directly-attached routing interfaces.
rip	Specifies that the advertisement filter is applied to FECs associated with RIP routes exported by OSPF.
static	Specifies that the advertisement filter is applied to FECs associated with static routes.
all	Specifies that unsolicited label mapping advertisements are originated for all routes of the specified type.
none	Specifies that no unsolicited label mapping advertisements are originated for the specified route type.
route-map	Specifies a route map is used to filter the origination of unsolicited label mapping advertisements for the specified route type.

Default

All — the default setting for the direct routing method.

None — the default setting for the RIP and static routing methods.

Usage Guidelines

Only the nlri-list route-map match operation keyword is supported for filtering origination of MPLS label advertisements.

You can configure how the advertisement filter is applied, as follows:

- `direct` — The advertisement filter is applied to the FECs associated with directly-attached routing interfaces.
- `rip` — The advertisement filter is applied to the FECs associated with RIP routes exported by OSPF.
- `static` — The advertisement filter is applied to the FECs associated with static routes.

You can configure the advertisement filter, as follows:

- `all` — All unsolicited label mappings are originated for all routes of the specified type (direct, RIP, or static). This is the default setting for direct routes.
- none — No unsolicited label mappings are originated for all routes of the specified type. This is the default setting for RIP and static routes.

- `route-map <route_map>` — The specified route map is used to permit or deny the origination of unsolicited label mappings for all routes of the specified type.

  The only supported route map match operation keyword is `nlri-list`. If selected, the `access_profile` parameter of the `nlri-list` keyword is compared to the FEC that is associated with each route.

NOTE

For more information on route maps, see the ExtremeWare Software Users Guide.

RIP routes are advertised with the Implicit NULL label and direct routes are advertised with an MPLS label, unless PHP is enabled.

Advertising labels for a large number of routes may increase the required number of labels that must be allocated by LSRs. Take care to ensure that the number of labels advertised by LERs does not overwhelm the label capacity of the LSRs.

Example

The following command configures a filter to be used by LDP when originating unsolicited label mapping advertisements for RIP routes:

```
config mpls ldp advertise rip all
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config mpls ldp advertise vlan

```
config mpls ldp advertise [add | delete] vlan <name>
```

Description

Configures LDP to originate an unsolicited label for the FECs associated with the directly attached routing interface of the specified VLAN.

Syntax Description

add	Originates an unsolicited label for the FECs associated with the directly attached routing interface of the specified VLAN
delete	Removes label origination of the direct route for the specified VLAN
vlan <name>	Specifies the name of the VLAN.

Default

N / A.

Usage Guidelines

Configures LDP to originate an unsolicited label for the FECs associated with the directly attached routing interface of the specified VLAN. The `delete` keyword removes label origination of the direct route for the specified VLAN. The LDP label origination configuration for directly attached routing interfaces can also be set using the `config mpls ldp advertise direct` command.

Example

The following command configures LDP to advertise a label for the direct route configured for VLAN vlan1:

```
config mpls advertise add vlan vlan1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config mpls php

```
config mpls php [enabled | disabled]
```

Description

Enables and disables penultimate hop popping (PHP) at the egress LSR. When enabled, PHP is requested on all LSPs for which the switch is the egress LSR.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

This command enables or disables whether PHP is requested by the egress LER.

When PHP is enabled, PHP is requested on all LSPs for which the switch is the egress LER.

PHP is requested by assigning the Implicit Null Label in an advertised mapping. PHP is always performed when requested by an egress LSR (for example, when the switch is acting as an intermediate LSR). The Implicit Null Label is always used in conjunction with routes exported by OSPF, regardless of the PHP configuration.

This command can only be executed when MPLS is disabled.

Example

The following command enables penultimate hop popping (PHP) at the egress LSR:

```
config mpls php enabled
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config mpls propagate-ip-ttl

```
config mpls propagate-ip-ttl [enabled | disabled]
```

Description

Enables or disables the propagation of the IP time-to-live (TTL) field for routed IP packets. When propagation is enabled, each LSR is viewed as a router hop from an IP TTL perspective. When propagation is disabled, the LSP is viewed as a point-to-point link between the ingress LSR and the egress LSR.

Syntax Description

This command has no arguments or variables.

Default

N / A.

Usage Guidelines

This command enables and disables the propagation of the IP TTL value for routed IP packets. The default setting is enabled.

NOTE

You must maintain identical `propagate-ip-ttl` *settings on all LERs in the MPLS domain. Not doing so may cause packets to loop endlessly and not be purged from the network if a routing loop is inadvertently introduced.*

When `propagate-ip-ttl` is disabled, the LSP is viewed as a point-to-point link between the ingress LSR and the egress LSR. Intermediate LSRs in the MPLS network are not viewed as router hops (from an IP TTL perspective). In this case, the IP TTL is decremented once by the ingress LSR and once by the egress LSR. When disabled, the MPLS TTL is set to 255 by the ingress LSR and is independent of the IP TTL.

When `propagate-ip-ttl` is enabled, each LSR is viewed as a router hop (from an IP TTL perspective). When a packet traverses an LSP, it emerges with the same TTL value that it would have had if it had traversed the same sequence of routers without being label-switched. When enabled, the MPLS TTL field is initially set to the IP TTL field at the ingress LSR, and the IP TTL field is set to the MPLS TTL by the egress LSR.

Example

The following command enables the propagation of the IP time-to-live (TTL) field for routed IP packets:

```
config mpls propagate-ip-ttl enabled
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config mpls qos-mapping

```
config mpls qos-mapping [dot1p-to-exp | exp-to-dot1p] [all |
<input_value>]/<output_value>
```

Description

Configures MPLS-specific QoS mappings.

Syntax Description

dot1p-to-exp	Specifies that mappings are used in performing the ingress LSR function. The value in this priority field is set based on the QoS classification performed by the ingress I/O module.
exp-to-dot1p	Specifies that mappings are used when performing label swapping as an intermediate LSR and when performing the egress LSR function.
all	Specifies to map all input values to the specified output value.
input_value	Specifies an input value.
output_value	Specifies an output value.

Default

Mapping tables are initialized such than an <input_value> of *n* is mapped to an <output_value> of *n*.

Usage Guidelines

The valid range of integers for the <input_value> and the <output_value> is 0 to 7. Two mappings are supported:

- dot1p-to-exp
- exp-to-dot1p

Dot1p-to-exp Mappings

The dot1p-to-exp mappings are used by the ingress LSR. When a non-MPLS ingress frame arrives at the MPLS module, the frame always contains an IEEE 802.1p priority field.

The value of the priority field is set based on the QoS classification performed by the ingress I/O module. The ingress I/O modules assign each packet to a hardware queue, based on the configured ExtremeWare QoS policies. There is a one-to-one mapping between the hardware queue and the 802.1p priority values that are inserted into frames forwarded to the MPLS module. For example, the 802.1p priority value is set to 0 for frames forwarded from hardware queue 0, set to 1 for frames forwarded from hardware queue 1, and so on.

The dot1p-to-exp table maps 802.1 priority values to MPLS EXP values. The table is completely flexible, such that any 802.1p priority <input_value> can be mapped to any EXP <output_value>. The EXP output_value is set in the MPLS header of the packet as it is forwarded to the MPLS network.

Exp-to-dot1p Mappings

The exp-to-dot1p mappings are used when the switch performs label swapping as an intermediate LSR and when the switch is the egress LSR. In both of these cases, the MPLS module receives an MPLS-encapsulated frame.

The EXP field in the frame is used as an <input_value> to the exp-to-dot1p table. The corresponding <output_value> is an 802.1p priority value. The 802.1p priority value is inserted into the frame before the frame is forwarded by the MPLS module.

The exp-to-dot1p table is completely flexible, such that any EXP <input_value> can be mapped to any 802.1p priority <output_value>.

The exp-to-dot1p table is also used by Packet over SONET (PoS) ports when classifying MPLS-encapsulated packets received from the SONET link. When a PoS port receives an MPLS-encapsulated packet from the SONET link, the packet is classified based on the EXP value in the MPLS shim header. The EXP value from the received frame is used as an index into the exp-to-dot1p mapping table to retrieve and 802.1p priority value. The frame is then assigned to a QoS profile, based on the retrieved 802.1p priority value. The mappings between 802.1p priority values and QoS profiles are configured using the following command:

```
config dot1p type
```

Example

The following command configures the dot1p-to-exp MPLS-specific QoS mappings:

```
config mpls qos-mapping dot1p-to-exp 0/1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config mpls rsvp-te add lsp

```
config mpls rsvp-te add lsp <lsp_name> path <path_name> {<profile_name>}
{primary | secondary}
```

Description

Adds an RSVP-TE LSP.

Syntax Description

lsp_name	Specifies the LSP name.
path_name	Specifies the path name
profile_name	Specifies the profile name.
primary	Specifies the primary LSP.
secondary	Specifies a secondary LSP.

Default

N / A.

Usage Guidelines

Both the <lsp_name> and <path_name> must be specified. The <lsp_name> parameter is a character string that is to be used to identify the LSP within the switch. The <lsp_name> string must begin with an alphabetic character and can contain up to 31 additional alphanumeric characters. The <profile_name> is optional. If omitted, the default profile is applied to the LSP. If no explicitly specified, the <path_name> defaults to the primary path. The LSP is immediately signaled as soon as it is configured. The maximum number of configurable LSPs is 1024.

Example

The following command adds a primary RSVP-TE LSP that takes the routed path named paththroughdenver:

```
config mpls rsvp-te add lsp lsptonyc path paththroughdenver
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config mpls rsvp-te add path

```
config mpls rsvp-te add path <path_name> [<ipaddress> | <host_name>] {from
<local_endpoint_vlan>}
```

Description

Adds a path to an RSVT-TE LSP.

Syntax Description

path_name	Specifies the path name.
ipaddress	Specifies the IP address.
hostname	Specifies the hostname.
local_endpoint_value	Specifies the local endpoint from which the path is signaled.

Default

N / A.

Usage Guidelines

The <path_name> and <ipaddress> or <host_name> must be specified for the path. The <path_name> parameter is a character string that is to used to identify the path within the switch. The <path_name> string must begin with an alphabetic character, and may contain up to 31 additional alphanumeric characters. Each <path_name> represents a routed path to a single IP destination.

If the <host_name> is specified, the DNS client on the switch must be configured so that the <host_name> can first be resolved to an IP address. Alternate routed paths to the same IP destination may be configured by adding additional <path_names> and specifying the same <ipaddress> or <host_name> as the path endpoint.

The RSVP-TE path is not signaled until an LSP is added with the specified <path_name>. If no explicit route objects are configured, the path will follow the best-routed path to the configured <ipaddress> (or IP address obtained from DNS name resolution). Optionally, the from keyword can be used to specify the <local_endpoint_vlan> from which the path is signaled. The maximum number of configurable paths is 255.

Example

The following command adds a path to 76.42.10.1 called paththroughdenver:

```
config mpls rsvp-te add path paththroughdenver 76.42.10.1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config mpls rsvp-te add profile

```
config mpls rsvp-te add profile <profile_name> {bandwidth <bps>}
{setup-priority <priority>} {hold-priority <priority>} {retry-timeout
<seconds>} {hop-count <number>} {ping-interval <seconds>} {metric [<metric>
| igp-tracking} {record [enabled | disabled]}
```

Description

Adds an RSVP-TE profile.

Syntax Description

profile_name	Specifies the profile name.
bandwidth	Specifies the reserved bandwidth for the LSP.
setup-priority	A value that is compared to the hold-priority of existing LSPs to determine if any of the existing LSPs need to be preempted to allow a higher priority LSP to be established.
hold-priority	A value that is compared to the setup-priority of existing LSPs to determine if any of the existing LSPs need to be preempted to allow a higher priority LSP to be established.
retry-timeout	Specifies the maximum number of seconds the switch allows for LSP setup.
ping-interval	Specifies how frequently an ICMP echo request is transmitted to the egress LSR IP address on the established LSP.
metric	Specifies a route metric used to determine if an established RSVP-TE LSP will actually be used to send data.
record	Specifies hop-by-hop path recording.

Default

N/A.

Usage Guidelines

A profile is a set of attributes that are applied to the LSP when the LSP is configured using the `config mpls rsvp-te add lsp` command. A default profile is provided which cannot be deleted, but can be applied to any configured LSP. The profile name for the default profile is *default*. The default profile parameter values are initially set to their respective default values. The maximum number of configurable profiles is 255 (one of which is reserved for the default profile).

The `bandwidth` parameter specifies the desired reserved bandwidth for the LSP. Any positive integer bps value is valid. Optionally, you can append the characters, k for kilobits, m for megabits, or g for gigabits, to the bps value to specify the unit of measure. If the k, m, or g, character is omitted, the unit of measure is assumed to be kilobits. The default bandwidth bps value is zero, which indicates that the QoS for the LSP is best effort. ExtremeWare does not support bandwidth reservation.

The `setup-priority` and `hold-priority` are optional parameters indicating the LSP priority. During path set up, if the requested bandwidth cannot be reserved through the LSR, the `setup-priority` parameter is compared to the `hold-priority` of existing LSPs to determine if any of the existing LSPs need to be preempted to allow a higher priority LSP to be established. Lower numerical values represent higher priorities. The `setup-priority` range is 0 to 7 and the default value is 7. The

hold-priority range is also 0 to 7 and is set equal to the setup-priority by default. ExtremeWare does not support LSP preemption.

The retry-timeout keyword specifies the maximum number of seconds the switch allows for LSP setup. If the LSP cannot be established within retry-timeout seconds, the LSP is resignaled. The default value for retry-timeout is 30 seconds with a configurable range of 5 to 600 seconds. The hop-count parameter limits the number of LSRs the path can traverse, including the ingress and egress router. The default hop-count value is 255 with a configurable range of two to 255.

After an LSP has established, the egress LSR may be optionally pinged to determine end-to-end path connectivity. If a ping response is not received within [2 * ping-interval – 1] seconds, the LSP is considered unavailable. The ping-interval keyword specifies how frequently an ICMP echo request is transmitted to the egress LSR IP address on the established LSP. The default ping-interval is zero, which indicates no end-to-end LSP health checking is performed. You can set the ping-interval value to any interval between 0 and 60 seconds.

The route metric is used to determine if an established RSVP-TE LSP will actually be used to send data. Whenever the configured metric is less than, or equal, to the calculated IGP metric, the LSP is used for sending routed IP traffic. In this case, the LSP is also used to send TLS data when the TLS tunnel is configured by specifying the tunnel LSP endpoint IP address. Traffic is distributed across up to four equal-cost LSPs. The valid metric values range from 1 to 65535. Specifying the igp-tracking keyword forces the route metric to track the underlying IGP metrics. If no IGP metric exists for the LSP (for example, the LSP traverses a RIP network), the metric is ignored. Tracking IGP metrics is the default behavior.

The record keyword is used to enable hop-by-hop path recording. The enabled keyword causes the record route object (RRO) to be inserted into the path message. The RRO is returned in the reserve message and contains a list of IPv4 subobjects that describe the RSVP-TE path. Path recording by default is disabled. When disabled, no RRO is inserted into the path message.

Example

The following command adds a profile with the configured attributes:

- Reserved bandwidth signaled is 100 Mbps
- Tunnel LSP setup priority is 1
- Tunnel LSP hold priority is 0
- Route recording is enabled

```
config mpls rsvp-te add profile customer1 bandwidth 100m setup-priority 1
hold-priority 0 record enabled
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config mpls rsvp-te delete lsp

```
config mpls rsvp-te delete lsp [<lsp_name> | all]
```

Description

Deletes an RSVP-TE LSP.

Syntax Description

lsp_name	Specifies the name of the LSP.

Default

N / A.

Usage Guidelines

Deleting an LSP name disassociates all configured paths with this LSP and all configuration information for the LSP name is deleted. LSPs cannot be deleted if the specified <lsp_name> has been configured as the LSP for a TLS tunnel. If you specify the all keyword, all LSPs not associated with a TLS tunnel are deleted.

Example

The following command deletes all RSVP-TE LSPs:

```
config mpls rsvp-te delete lsp all
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config mpls rsvp-te delete path

```
config mpls rsvp-te delete path [<path_name> | all]
```

Description

Deletes an RSVP-TE path.

Syntax Description

path_name	Specifies the name of the path.

Default

N / A.

Usage Guidelines

This command deletes a configured MPLS RSVP-TE routed path with the specified <path_name>. All associated configuration information for <path_name> is deleted. A path cannot be deleted as long as the <path_name> is associated with an LSP. If the all keyword is specified, all paths not associated with an LSP are deleted.

Example

The following command deletes all RSVP-TE paths:

```
config mpls rsvp-te delete path all
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config mpls rsvp-te delete profile

```
config mpls rsvp-te delete profile [<profile_name> | all]
```

Description

Deletes an RSVP-TE path profile.

Syntax Description

profile_name	Specifies the name of the profile.

Default

N / A.

Usage Guidelines

This command deletes a configured RSVP-TE profile with the specified profile name. The default profile cannot be deleted. If a profile is associated with a configured LSP, the profile cannot be deleted. If you specify the `all` keyword, all profiles not associated with an LSP are deleted (except for the default profile).

Example

The following command deletes all RSVP-TE path profiles:

```
config mpls rsvp-te delete profile all
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config mpls rsvp-te lsp add path

```
config mpls rsvp-te lsp <lsp_name> add path <path_name> {<profile_name>}
{secondary | primary}
```

Description

Adds a path to an RSVP-TE LSP.

Syntax Description

lsp_name	Specifies the name of a configured LSP.
path_name	Specifies the path name.
profile_name	Specifies the profile name.
primary	Specifies the primary path.
secondary	Specifies a secondary path.

Default

N/A.

Usage Guidelines

The <lsp_name> must represent a configured LSP. Only one primary path and up to two secondary paths can be added per <lsp_name>. The <path_name> specified defaults to primary when no primary path has been configured for <lsp_name> and defaults to secondary if the primary path has been previously configured for <lsp_name>.

You do not need to configure the primary path for an LSP. Each <path_name> added to an <lsp_name> must be unique, but a <path_name> can be associated with multiple LSP names.

All configured primary and secondary paths for the <lsp_name> must have the same endpoint IP address. For example, three paths can be configured for the <lsp_name>, but all paths should represent different topological paths through the network to the same LSP endpoint.

Adding a secondary <path_name> designates a path as a hot-standby redundant path, used in the event that the primary or secondary path cannot be established or fails. Provided the <path_name> has not already been established, all path names are signaled as soon as they are associated with an <lsp_name>. If the primary <path_name> fails, is not configured, or cannot be established after the specified LSP retry-timeout, one of the configured secondary paths may become the active path for <lsp_name>. All of the secondary paths have equal preference; the first one available is chosen. If at any time the primary path is established, <lsp_name> immediately switches to using the primary path. If a secondary path fails while in use, the remaining configured secondary paths can become the active path for <lsp_name>.

Example

The following command adds a secondary path named paththroughdc for the specified LSP:

```
config mpls rsvp-te lsp lsptonyc add path paththroughdc secondary
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config mpls rsvp-te delete path

```
config mpls rsvp-te delete path [<path_name> | all]
```

Description

Deletes an RSVP-TE path.

Syntax Description

path_name	Specifies the name of the path.

Default

N/A.

Usage Guidelines

This command deletes a configured MPLS RSVP-TE routed path with the specified <path_name>. All associated configuration information for <path_name> is deleted. A path cannot be deleted as long as the <path_name> is associated with an LSP. If the all keyword is specified, all paths not associated with an LSP are deleted.

Example

The following command deletes all RSVP-TE paths.

```
config mpls rsvp-te delete path all
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config mpls rsvp-te add ero

```
config mpls rsvp-te path <path_name> add ero [ipaddress
<ipaddress/masklength> | <host_name>] {strict | loose} {order <number>}
```

Description

Adds an RSVP-TE explicit route.

Syntax Description

path_name	Specifies the path name.
ipaddress/masklength	Specifies an LSR using either a /32 address, which may represent an LSR router ID, loopback address, or direct router interface, or an IP prefix, which represents a directly connected subnet.
strict	Specifies a strict subobject.
loose	Specifies a loose subobject.
order <number>	Specifies the LSR path order.

Default

N/A.

Usage Guidelines

This command adds an IP address to the explicit route object (ERO) for the specified path name. The RSVP-TE routed path may be described by a configured sequence of the LSRs and/or subnets traversed by the path. Each defined LSR or subnet represents an ERO subobject. Up to 64 subobjects can be added to each path name.

When specifying an LSR using the <host_name> parameter, the DNS client on the switch must be configured so that the <host_name> can first be resolved to an IP address. The ipaddress keyword identifies an LSR using either a /32 address, which may represent an LSR router ID, loopback address, or direct router interface, or an IP prefix, which represents a directly connected subnet. Each IP address or prefix is included in the ERO as an IPv4 subobject. Each specified subobject must be topologically adjacent to the next subobject, as listed in the ERO. If the subobject matches a direct router interface or a directly attached subnet, the switch verifies that the path message is received on the matching router interface. If the LSR specified matches the OSPF router ID or a configured loopback IP address, the router interface on which the packet is received is ignored.

If the IP address is specified as strict, the strict subobject must be topologically[1] adjacent to the previous subobject as listed in the ERO. If the IP address is specified as loose, the loose subobject is not required to be topologically adjacent to the previous subobject as listed in the ERO. If omitted, the default subobject attribute is strict. Each IP address or prefix is included in the ERO as an IPv4 subobject.

If the subobject matches a direct router interface or a directly attached subnet, the switch verifies that the path message is received on the matching router interface. If the LSR specified matches the OSPF

1. The LSP next hop matches either the interface IP address or the OSPF router-id of the immediate neighbor LSR.

router ID or a configured loopback IP address, the router interface which the packet is received is ignored.

The LSR path order is optionally specified using the `order` keyword. The `order number` parameter is an integer value from 1 to 65535. IP prefixes with a lower number are sequenced before IP prefixes with a higher number. You can specify multiple paths and assign them an order number. The order number determines the path that the LSP follows. Thus, the LSP path follows the configured path of the IP prefix with the order value from low to high. If the `order` keyword is not specified, the number value for the LSR defaults to a value 100 higher than the current highest number value.

If the list of IP prefixes, added to the path, does not reflect an actual path through the network topology, the path message is returned with an error from a downstream LSR and the LSP is not established.

The order of a configured subobject can not be changed. The ERO subobject must be deleted and re-added using a different order. If a subobject is added to or deleted from the ERO while the associated LSP is established, the path is torn down and is resignaled using the new ERO.

Duplicate ERO subobjects are not allowed. Defining an ERO for the path is optional. If you do not configure an ERO, the path is signaled along the best-routed path and the ERO is not included in the path message. When the last subobject in the ERO of the path message is reached and the egress IP node of the path has not been reached, the remaining path to the egress node is signaled along the best-routed path. Specification of an ERO could lead to undesirable routed paths, so you should be careful when terminating the ERO routed-path definition prior to the configured path egress node.

Example

The following command adds a strict ERO subobject of 192.18.32.5 to the specified path.

```
config mpls rsvp-te path paththroughdenver add ero ipaddress 192.18.32.5
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config mpls rsvp-te delete ero

```
config mpls rsvp-te path <path_name> delete ero [all | ipaddress
<ipaddress/masklength> | <host_name> | order <number>]
```

Description

Deletes an RSVP-TE explicit route.

Syntax Description

path_name	Specifies the path name.
ipaddress/masklength	Specifies an LSR using either a /32 address, which may represent an LSR router ID, loopback address, or direct router interface, or an IP prefix, which represents a directly connected subnet.
strict	Specifies a strict subobject.
loose	Specifies a loose subobject.
order <number>	Specifies the LSR path order.

Default

N / A.

Usage Guidelines

This command deletes an LSR or subnet from the ERO for the specified path name. The LSR is specified using the `ipaddress`, `<host_name>`, or `order` parameter. If an LSR is deleted from an ERO while the associated LSP is established, the path is torn down and is resignaled using a new ERO. Use the `all` keyword to delete the entire ERO from the path name. When there is no configured ERO, the path is no longer required to take an explicit routed path. The path is then signaled along the best-routed path and no ERO is included in the path message.

Example

The following command deletes all configured ERO subobjects from the specified path:

```
config mpls rsvp-te path paththroughdc delete ero all
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config mpls rsvp-te profile

```
config mpls rsvp-te profile <profile_name> {bandwidth <bps>} {hop-count
<number>} {setup-priority <priority>} {hold-priority <priority>}
{retry-timeout <seconds>} {ping-interval <seconds>} {metric [<metric> |
igp-tracking]} {record [enabled | disabled]}
```

Description

Configures an existing RSVP-TE profile.

Syntax Description

profile_name	Specifies the profile name.
bandwidth	Specifies the reserved bandwidth for the LSP.
setup-priority	A value that is compared to the hold-priority of existing LSPs to determine if any of the existing LSPs need to be preempted to allow a higher priority LSP to be established.
hold-priority	A value that is compared to the setup-priority of existing LSPs to determine if any of the existing LSPs need to be preempted to allow a higher priority LSP to be established.
retry-timeout	Specifies the maximum number of seconds the switch allows for LSP setup.
ping-interval	Specifies how frequently an ICMP echo request is transmitted to the egress LSR IP address on the established LSP.
metric	Specifies a route metric used to determine if an established RSVP-TE LSP will actually be used to send data.
record	Specifies hop-by-hop path recording.

Default

N/A.

Usage Guidelines

This command configures RSVP-TE attributes for the specified profile. The `<profile_name>` must have been previously added. All of the LSP profile values are updated dynamically. For LSPs configured with this profile, the LSP parameters are updated automatically with the sending of the next refresh path message. If the metric is changed, all LSPs using this profile are rechecked against the calculated IGP metric. In some cases, the LSP may be torn down because of a profile configuration change. For example, if the bandwidth value is increased, the LSRs along the existing path may not be able to accommodate the additional reserved bandwidth. In this scenario, the LSP is torn down and resignaled.

Example

The following command configures the attributes for the specified profile:

```
config mpls rsvp-te profile customer1 ping-interval 2
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config mpls rsvp-te vlan

```
config mpls rsvp-te vlan [<name> | all] {hello-interval <seconds>}
{refresh-time <seconds>} {summary-refresh-time <seconds>} {bundle-time
<seconds>} {keep-multiplier <number>}
```

Description

Configures RSVP-TE protocol parameters

Syntax Description

name	Specifies the VLAN name.
hello-interval	Specifies the RSVP hello packet transmission interval.
refresh-time	Specifies the interval for sending refresh path messages.
bundle-time	Specified the maximum amount of time a transmit buffer is held so that multiple RSVP messages can be bundled into a single PDU.
summary-refresh-time	Specifies the time interval for sending summary refresh RSVP messages.

Default

N / A.

Usage Guidelines

This command configures the RSVP-TE protocol parameters for the specified VLAN. The RSVP-TE keyword `all` indicates that the configuration changes apply to all RSVP-TE enabled VLANs.

The `hello-interval` time specifies the RSVP hello packet transmission interval. The RSVP hello packet is used by the switch to detect when a RSVP-TE peer is no longer reachable. If an RSVP hello packet is not received from a peer with [hello-interval * `keep-multiplier`] seconds, the peer is declared down and all RSVP sessions to and from that peer are torn down. The default `hello-interval` time is three seconds with a valid range from one to 60 seconds.

The `refresh-time` specifies the interval for sending refresh path messages. RSVP refresh messages provide “soft state” link-level keep-alive information for previously established paths and enables the switch to detect when an LSP is no longer active. RSVP sessions are torn down if an RSVP refresh message is not received from a neighbor within [(`keep-multiplier` + 0.5) * 1.5 * `refresh-time`] seconds. The default `refresh-time` is 30 seconds and the default `keep-multiplier` value is three. The minimum and maximum `refresh-time` values are one and 36,000 seconds (or one hour) respectively. The minimum and maximum `keep-multiplier` values are one and 255 respectively.

The `bundle-time`, specified in tenths of a second, indicates the maximum amount of time a transmit buffer is held so that multiple RSVP messages can be bundled into a single PDU. The default `bundle-time` is zero, indicating that RSVP message bundling is not enabled. The `bundle-time` value may be set to any value between zero and 30 (or 3 seconds).

The `summary-refresh-time`, specified in tenths of a second, indicates the time interval for sending summary refresh RSVP messages. The `summary-refresh-time` must be less than the configured `refresh-time`. The default `summary-refresh-time` is zero, indicating that no summary refresh RSVP

messages are sent. The summary-refresh-time value may be set to any value between zero to 100 (or 10 seconds).

If configured, the bundled and summary refresh RSVP messages are only sent to RSVP-TE peers supporting RSVP refresh reduction.

Example

The following command configures the rsvp-te interface parameters for VLAN vlan1.

```
config mpls rsvp-te vlan vlan1 hello-interval 2 refresh-time 5
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config mpls tls-tunnel vlan mode

```
config mpls tls-tunnel vlan [<name>] mode [hub | mesh]
```

Description

Configures the VPN flood mode.

Syntax Description

name	Specifies the name of the VLAN.
hub	Specifies hub as the flood mode.
mesh	Specifies mesh as the flood mode.

Default

N / A.

Usage Guidelines

This command configures the broadcast and unknown packet-forwarding behavior for the specified TLS VLAN. The TLS VPN flood mode options are hub and mesh. When two or more TLS tunnels are configured for the same TLS VLAN, each configured TLS tunnel and the local TLS VLAN are treated as separate bridge ports within a single layer 2 broadcast domain.

When the mode is configured as hub, the TLS LSR behavior is similar to a repeater. All received broadcast and unknown unicast packets are flooded out every port, except for the port on which the packet was received. When the mode is configured as mesh, the TLS LSR only floods packets received from the local TLS VLAN for transmission onto every TLS tunnel. Traffic received from a TLS tunnel is forwarded only to the local TLS VLAN. The default mode is mesh.

Example

The following command configures the flood mode for VLAN vlan2 as mesh:

```
config mpls tls-tunnel vlan vlan2 mode mesh
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config mpls vlan ip-mtu

```
config mpls vlan [<name> | all] ip-mtu <number>
```

Description

Configures the IP MTU for frames transmitted onto MPLS LSPs via the specified egress VLAN. The range is 46 to 9216 (using jumbo frame sizes).

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.
number	Specifies an IP MTU size.

Default

1500 bytes.

Usage Guidelines

This command configures the IP MTU for frames transmitted onto MPLS LSPs via the specified egress VLAN. The default settings is 1500 bytes. If `all` is selected, the configuring MTU applies to all MPLS-enabled VLANs.

This command applies to the ingress LSR only when a received IP packet is destined for an MPLS LSP. In this case, if the length of the IP packet exceeds the configured MTU size for the egress VLAN and the Don't Fragment (DF) bit is *not* set in the IP header of the packet, the packet is fragmented before it is forwarded onto an MPLS LSP. If the DF bit is set in the packet header, Path MTU Discovery starts.

Fragmentation is based on either the minimum value of the configured MPLS IP MTU size or the configured IP MTU size for the egress VLAN. (The IP MTU size is configured using the `config ip-mtu <number> vlan <name>` command.)

Configure the MPLS IP MTU so that the addition of the MPLS label stack the link layer header does not cause the packet to be too large to be transmitted on the egress ports. To avoid potential problems, enable jumbo frame support on all ports that are members of an MPLS VLAN.

Example

The following command configures the IP MTU for frames transmitted onto MPLS LSPs:

```
config mpls vlan vlan1 ip-mtu 1550
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

config mpls vlan ldp propagate

```
config mpls vlan [<name> | all] ldp propagate [all | none | route-map
<route_map>]
```

Description

Configures a filter to be used by LDP when propagating unsolicited label mappings to all LDP neighbors on one or all VLANs.

Syntax Description

name	Specifies a VLAN name.
all	Specifies all VLANs.
all	Specifies that all unsolicited label mappings are propagated to the VLAN.
none	Specifies that no unsolicited label mappings are propagated to the VLAN.
route_map	Specifies the route map used to permit or deny the propagation of unsolicited label mappings to the VLAN.

Default

All unsolicited label mappings are propagated to the VLAN.

Usage Guidelines

This command configures a filter to be used by LDP when propagating unsolicited label mappings to all LDP neighbors on the specified VLAN. If all VLANs are selected, the settings of this command apply to all MPLS-enabled VLANs.

Example

The following command configures a filter to be used by LDP when propagating unsolicited label mappings to *vlan1*:

```
config mpls vlan vlan1 ldp propagate route-map bgp_out
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

disable mpls

```
disable mpls
```

Description

Disables MPLS on the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Disabling MPLS causes all LSPs to be released and all LDP neighbor sessions to be terminated.

Example

The following command globally disables MPLS on the switch:

```
disable mpls
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

disable ospf originate-router-id

```
disable ospf originate-router-id
```

Description

Disables distribution of a route for the OSPF router ID in the router LSA.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

When this function is enabled, OSPF includes a link with the router ID IP address and a mask of 255.255.255.255 in the router LSA. The link type is stub and the metric is 0.

When disabled, OSPF does not include a link with the router ID IP address in the router LSA

Example

The following command disables the distribution of a route for the OSPF router ID in the router LSA:

```
disable ospf originate-router-id
```

History

This command was available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

enable mpls

```
enable mpls
```

Description

Enables MPLS on the switch.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

Disabling MPLS causes all LSPs to be released and all LDP neighbor sessions to be terminated.

Example

The following command globally enables MPLS on the switch:

```
enable mpls
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

enable ospf originate-router-id

```
enable ospf originate-router-id
```

Description

Enables distribution of a route for the OSPF router ID in the router LSA.

Syntax Description

This command has no arguments or variables.

Default

Disabled.

Usage Guidelines

When this function is enabled, OSPF includes a link with the router ID IP address and a mask of 255.255.255.255 in the router LSA. The link type is stub and the metric is 0.

When disabled, OSPF does not include a link with the router ID IP address in the router LSA

Example

The following command enables the distribution of a route for the OSPF router ID in the router LSA:

```
enable ospf originate-router-id
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

show mpls

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

Description

Displays MPLS configuration information for one or all VLANs. Omitting the vlan keyword displays information for all VLANs.

Syntax Description

name	Specifies a VLAN name.
detail	Specifies to display the information in detailed format.

Default

N / A.

Usage Guidelines

When the `vlan` parameter is omitted, this command displays the values of all MPLS configuration parameters that apply to the entire switch, the current status of peer LSRs, and a list of the VLANs for which MPLS is enabled.

When the `vlan` parameter is specified, this command displays the current values of the MPLS configuration parameters that are specific to the VLAN.

If the optional `detail` keyword is specified, additional detailed VLAN information is displayed.

Example

The following command displays MPLS configuration information for the VLAN *accounting*:

```
show mpls vlan accounting
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

show mpls forwarding

```
show mpls forwarding {summary | detail | inactive | host <ipaddress>
{detail | inactive} | prefix <ipaddress/masklength> {detail | inactive} |
rsvp-te <ipaddress> {detail}}
```

Description

Displays information from the FEC-to-NHLFE database, used when forwarding non-MPLS packets onto an LSP. Also displays information for RSVP-TE LSPs.

Syntax Description

summary	Displays only the summary route information associated with labeled paths.
host	Displays information for a single FEC.
prefix	Displays information for a single FEC.
rsvp-te	Displays only the RSVP-TE forwarding label mapping
inactive	Causes inactive mappings to be displayed. This keyword does not apply to the `rsvp-te` keyword, because RSVP-TE operates in DoD mode.

Default

N/A.

Usage Guidelines

This command displays information from the Forwarding Equivalence Class (FEC)-to-Next Hop Label Forwarding Entry (NHLFE) database. This command also displays information for RSVP-TE LSPs.

If the `host` or `prefix` keywords are specified, summary information is displayed for a single FEC. Use the `summary` keyword to display summary route information associated with labeled paths.

By default, the information displayed includes:

- Next hop IP address
- Outgoing label
- Interface number of the outgoing VLAN

If the `detail` keyword is specified, the following additional information is displayed:

- Outgoing port number
- Counts of packets and bytes that have been transmitted using the database entry

By default, information is displayed for active mappings. To display information for liberally-retained inactive mappings, use the inactive keyword. An inactive mapping is a mapping that was received from an LDP peer, but is not being used to reach the associated FEC. Using the `inactive` keyword causes inactive mappings to be displayed. The inactive keyword does not apply to RSVP-TE LSPs, because RSVP-TE operates in downstream-on-demand mode.

Example

The following command displays information from the FEC-to-NHLFE database:

```
show mpls forwarding prefix 10.1.1.1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

show mpls interface

```
show mpls interface {ldp | targeted-ldp | rsvp-te}
```

Description

Displays targeted LDP and RSVP-TE interface information.

Syntax Description

ldp	Specifies LDP interfaces.
targeted-ldp	Specifies targeted LDP interfaces.
RSVP-TE	Specifies RSVP-TE interfaces.

Default

N / A.

Usage Guidelines

Displays targeted LDP and RSVP-TE interface information, including targeted LDP and RSVP-TE peer IP address and peer state. Specifying the keyword `ldp`, `targeted-ldp`, or `rsvp-te` limits the information displayed to only those interface types.

Example

The following command displays interface information for RSVP-TE interfaces:

```
show mpls interface rsvp-te
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

show mpls label

```
show mpls label {summary | detail | <label_number> {detail} | host
<ipaddress> {detail} | prefix <ipaddress/masklength> {detail} | rsvp-te
<ipaddress> {detail}}
```

Description

Displays information from the Incoming Label Map (ILM), used when forwarding packets that arrive as labeled MPLS packets.

Syntax Description

summary	Specifies the number of labels allocated form each label range partition.
detail	Specifies to display the information in detailed format.
label_number	Specifies an MPLS label number.
host <ipaddress>	Specifies a particular host FEC type.
prefix	Specifies a particular prefix FEC type.
rsvp-te	Specifies only RSVP-TE assigned labels

Default

N / A.

Usage Guidelines

This command displays information from the Incoming Label Map (ILM), which is used when forwarding packets that arrive labeled as MPLS packets.

When the `label_number` parameter is omitted, summary information is displayed for all incoming label assignments that have been made by the switch. When the `label_number` is specified, summary information is displayed for the label.

Use the `fec` keyword to display the label associated with an FEC. You can specify both host and prefix FEC types. The `summary` keyword displays the number of labels allocated from each label range partition.

By default, the information displayed includes:

- Next hop IP address
- Outgoing and incoming labels
- Interface number of the outgoing VLAN
- FEC associated with the incoming label

If the detail keyword is specified, the following additional information is displayed:

- Outgoing port number
- Counts of packets and bytes that have been received with the incoming label

- Counts of packets and bytes that have been transmitted with the outgoing label
- LSP type

This command also displays information from the Incoming Label Map (ILM) for RSVP-TE LSPs.

Example

The following command displays the summary information from the Incoming Label Map:

```
show mpls label summary
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

show mpls ldp

```
show mpls ldp {<ipaddress>} {detail}
```

Description

Displays MPLS LDP session information for one or all LSP sessions.

Syntax Description

ipaddress	Specifies an IP address.
detail	Specifies to display the information in detailed format.

Default

N / A.

Usage Guidelines

Omitting the ipaddress parameter displays LDP session information for all LDP sessions.

This command displays information about the status of LDP peers. Summary information is displayed for all known LDP peers and LDP peer sessions. If you specify the `<ipaddress>` of the LDP peer, information for a single LDP peer is displayed. To display additional information in the comprehensive detailed format, use the `detail` keyword.

Displayed summary information includes:

- Peer type (targeted or not targeted)
- Peer status
- Peer sessions
- Peer session state

If you specify the `detail` keyword, the following additional information is displayed:

- LDP error counts
- LDP status timers
- Maximum PDU length

Example

The following command displays MPLS LDP session information for the LDP entity 10.1.1.1:

```
show mpls ldp 10.1.1.1
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

show mpls qos-mapping

```
show mpls qos-mappings
```

Description

Displays MPLS-specified QoS mappings for dot1p-to-exp and exp-to-dot1p.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Configured mappings for both dot1p-to-exp and exp-to-dot1p are displayed.

Example

The following command displays MPLS QoS mapping information:

```
show mpls qos-mappings
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

show mpls rsvp-te

```
show mpls rsvp-te {<ipaddress>} {detail}
```

Description

Displays RSVP-TE LSP configuration information.

Syntax Description

ipaddress	Specifies the IP address of the RSVP-TE interface.
detail	Specifies to display the information in detailed format.

Default

N / A.

Usage Guidelines

This command displays information about the status of RSVP-TE enabled interfaces. Summary information is displayed for all known RSVP-TE peers including the peer IP address and peer status. If you specify the `ipaddress` of the RSVP-TE interface, the information for a single RSVP-TE interface is displayed. Additional information is displayed in the detailed format if you specify the optional `detail` keyword. The more detailed RSVP-TE information includes the number and type of RSVP messages transmitted through the local RSVP-TE interface.

Example

The following displays detailed information about all configured RSVP-TE LSPs:

```
show mpls rsvp-te detail
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

show mpls rsvp-te lsp

```
show mpls rsvp-te lsp {<lsp_name>} {detail}
```

Description

Displays the RSVP-TE LSP.

Syntax Description

lsp_name	Specifies the name of the LSP.
detail	Specifies to display the information in detailed format.

Default

N / A.

Usage Guidelines

This command displays the configuration and status information for RSVP-TE LSPs. Information is listed in tabular format and includes the LSP name, LSP state, active path name, bandwidth requested, bandwidth actually reserved, ERO flag, egress LSR, LSP up-time, and RSVP error codes (if LSP setup failed). If you specify a specific LSP name, only information for the specified LSP is displayed. If you specify the optional `detail` keyword, additional information is displayed for each LSP. The detailed information includes a list of all configured paths, including the path state, error codes for the LSP associated with each path, up-time for each LSP, the bound profile name, and a list of TLS tunnels configured to use the LSP.

Example

The following displays the configuration and status information for all configured RSVP-TE LSPs in detailed format:

```
show mpls rsvp-te lsp detail
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

show mpls rsvp-te path

```
show mpls rsvp-te path {<path_name>} {detail}
```

Description

Displays the RSVP-TE routed path.

Syntax Description

path_name	Specifies the name of the path.
detail	Specifies to display the information in detailed format.

Default

N / A.

Usage Guidelines

This command displays the configuration and status information for MPLS RSVP-TE routed paths. Information is listed in tabular format and includes the path name, path endpoint LSR IP address, and local VLAN (if configured). If the path endpoint is specified as a host name, the host name and the DNS resolved IP address are both displayed. If a specific path name is specified, only information for the specified path is displayed. If you specify the optional `detail` keyword, the list of subobjects specified for the explicit route object and any LSPs that are configured to use the path are displayed.

Example

The following displays information about all RSVP-TE routed paths in detailed format:

```
show mpls rsvp-te path detail
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

show mpls rsvp-te profile

```
show mpls rsvp-te profile {<profile_name>}
```

Description

Displays the RSVP-TE path profile.

Syntax Description

profile_name	Specifies the name of the profile.

Default

N / A.

Usage Guidelines

By default, this command displays all configured profile parameters for the specified profile. If the profile name is omitted, the profile parameter values for all configured LSP profiles are displayed.

Example

The following command displays the profile parameter values for all configured LSP profiles:

```
show mpls rsvp-te profile
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

show mpls tls-tunnel

```
show mpls tls-tunnel {summary | detail | <tunnel_name> {detail} | vlan
<vlan_name> {detail}}
```

Description

Displays configuration and status information for TLS tunnels.

Syntax Description

summary	Specifies to display summary TLS tunnel counts.
detail	Specifies to display the information in detailed format.
tunnel_name	Specifies a TLS tunnel name.
vlan_name	Specifies a VLAN name.

Default

N / A.

Usage Guidelines

This command displays configuration and status information for one or all TLS tunnels. The information displayed for each tunnel includes:

- The values of all configuration parameters for the tunnel.
- The current status of the tunnel LSP.
- Transmit and receive counts in terms of packets and bytes.

If the optional `detail` keyword is specified, TLS tunnel information is displayed using the comprehensive detail format.

If the optional `summary` keyword is specified, summary TLS tunnel counts are displayed. The summary counters displayed include the total number of active static and dynamic TLS tunnels.

Example

The following command displays configuration and status information for the TLS tunnel *rt40*:

```
show mpls tls-tunnel rt40
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

unconfig mpls

```
unconfig mpls
```

Description

Resets MPLS configuration parameters to the default settings.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command resets the following configuration parameters:

- IP-MTU
- LDP propagation filter settings on all VLANs
- LDP advertisement filter settings
- LDP session timers
- RSVP-TE interface parameters
- RSVP-TE profile parameters
- Settings for propagate-ip-ttl
- QoS mapping tables

Example

The following command resets MPLS configuration parameters to the default settings:

```
unconfig mpls
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

unconfig mpls

```
unconfig mpls [hello-hold-time | session-keep-alive-time]
```

Description

Restores the default values for hello-hold-time or session-keep-alive-time.

Syntax Description

hello-hold-time	Specifies a hello hold time.
session-keep-alive-time	Specifies a session keep alive time.

Default

The default hello-hold-time is 15 seconds.

The default session-keep-alive-time is 40 seconds.

Usage Guidelines

This command can only be executed when MPLS is disabled.

The hello-hold-time is the amount of time, in seconds, an LSR maintains a record of the label space requested by potential LDP peers. An LSR must receive an LDP hello packet at least hello-hold-time seconds after the last hello packet was received, or the LSR concludes that the LDP peer has failed or no longer wishes to label switch using the previously advertised label space.

The session-keep-alive-time specifies the minimum amount of time, in seconds, that an LSR must receive an LDP PDU from an LDP peer to which it has an established LDP session. If an LDP PDU is not received within the specified session-keep-alive-time since the reception of the last LDP PDU, the LDP session is torn down.

Example

The following command restores the default values for hello-hold-time:

```
unconfig mpls hello-hold-time
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

unconfig mpls qos-mapping

```
unconfig mpls qos-mapping [dotp-to-exp | exp-to-dot1p | lsp <lsp_name>]
```

Description

Restores the default values for the specified QoS mapping table.

Syntax Description

dot1p-to-exp	Specifies dot1p-to-exp mapping.
exp-to-dot1p	Specifies exp-to-dot1p mapping.
lsp_name	Specifies the name of an LSP.

Default

N / A.

Usage Guidelines

The default contents of either QoS mapping table maps an input value of *n* to an output value of *n*.

Example

The following command restores the default values for the dot1p-to-exp QoS mapping table:

```
unconfig mpls qos-mapping dot1p-to-exp
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on v6.1.8b12.

Platform Availability

This command is available on the BlackDiamond 6800 series chassis-based system only.

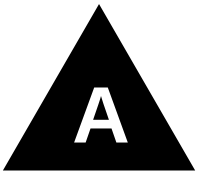

Configuration and Image Commands

This appendix describes the following commands:

- Commands related to downloading and using a new switch software image
- Commands related to saving, uploading, and downloading switch configuration information
- Commands related to the BootROM and switch rebooting

The switch software *image* contains the executable code that runs on the switch. An image comes preinstalled from the factory. The image can be upgraded by downloading a new version from a Trivial File Transfer Protocol (TFTP) server on the network.

A switch can store up to two images; a primary and a secondary image. You can download a new image into either one of these, and you can select which image will load on the next switch reboot.

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. To retain the settings, and have them load when you reboot the switch, you must save the configuration to nonvolatile storage.

A switch can store two different configurations: a primary and a secondary configuration. You can select to which configuration you want the changes saved, and which configuration will be used on the next switch reboot.

The BootROM initializes certain important switch variables during the switch boot process. In specific situations, the BootROM can be upgraded by download from a TFTP server on the network.

config download server

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

Description

Configures the TFTP server(s) used by a scheduled incremental configuration download.

Syntax Description

primary	Specifies that the following parameters refer to the primary TFTP server.
secondary	Specifies that the following parameters refer to the secondary TFTP server.
ip address	Specifies the IP address of the TFTP server from which the configuration should be obtained.
hostname	Specifies the hostname of the TFTP server from which the configuration should be obtained.
filename	Specifies the filename on the server that contains the configuration to be downloaded.

Default

N/A.

Usage Guidelines

This command must be executed before scheduled configuration downloads can be performed.

Use of the <hostname> parameter requires that DNS be enabled.

Example

The following command specifies that scheduled incremental downloads into the primary configuration space be done from the server named *tftphost*, from the ASCII file *primeconfig.txt* (residing in directory *\configs\archive* on the server).

```
config download server primary tftphost \configs\archive\prime_config.txt
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

download bootrom

```
download bootrom [<ip address> | <hostname>] <filename> {slot <slot>}
```

Description

Downloads a BootROM image from a TFTP server after the switch has booted. The downloaded image replaces the BootROM in the onboard FLASH memory.

Syntax Description

ip address	Specifies the IP address of the TFTP server.
hostname	Specifies the hostname of the TFTP server.
filename	Specifies name of the file on the server that contains the bootROM image.
slot	Specifies the slot where a PoS or MPLS module is installed.

Default

N/A.

Usage Guidelines

Upgrade the BootROM only when asked to do so by an Extreme Networks technical representative.

If this command does not complete successfully it could prevent the switch from booting. In the event the switch does not boot properly, some boot option functions can be accessed through a special BootROM menu (see the ExtremeWare Software User Guide).

Use of the <hostname> parameter requires that DNS be enabled.

Example

The following command downloads a bootROM image from the tftp server *tftphost* from the file *bootimages* (residing in directory *\images* on the server):

```
download bootrom tftphost \images\bootimage
```

History

This command was first available in ExtremeWare 4.0.

This command was modified in the ExtremeWare IP Services Technology Release based on 6.1.8b12 to support download to a PoS or MPLS module.

Platform Availability

This command is available on all platforms.

download configuration

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

Description

Downloads a previously saved ASCII configuration file from a specific TFTP server host.

Syntax Description

ip address	Specifies the IP address of the TFTP server from which the configuration should be obtained.
hostname	Specifies the hostname of the TFTP server from which the configuration should be obtained.
filename	Specifies the path and filename of a saved ASCII configuration.
incremental	Specifies an incremental configuration download (v 6.0 or later).

Default

N/A.

Usage Guidelines

Unless you specify the incremental keyword, this command does a complete download, resetting the current switch configuration and replacing it with the new downloaded configuration. You will be prompted to reboot the switch after the download is complete.

Use the `incremental` keyword to specify an incremental or partial configuration download. In this case, the commands specified in the incremental download file are executed, but configuration settings not specified in the file are left intact. No reboot is required.

The new configuration information is stored in switch runtime memory, and is not retained if the switch has a power failure. After the switch has rebooted, you should save the configuration to the primary or secondary configuration area to retain it through a power cycle. You can include a save command at the end of the configuration file to have the save done at the end of the download.

The file on the server is assumed to be located relative to the TFTP server base directory. You can specify a path as part of the file name.

Use of the <hostname> parameter requires that DNS be enabled.

Example

The following command clears the current switch configuration, and downloads a new full configuration from the tftp server *tftphost*. It uses the configuration from the file *stdconfigs.txt* residing in the subdirectory *configs\archive* of the TFTP server base directory on the server:

```
download configuration tftphost configs\archive\stdconfig.txt
```

The following command downloads a partial configuration from the tftp server *tftphost* from the file *modifyconfig.txt* (residing in the subdirectory *configs\archive* on the server):

```
download configuration tftphost configs\archive\modifyconfig.txt incremental
```

History

This command was first available in ExtremeWare 2.0.

Support for the <hostname> parameter was introduced in ExtremeWare 4.0.

Support for incremental downloads was introduced in ExtremeWare 6.0.

Platform Availability

This command is available on all platforms. The incremental download option is available on the *"i"* series platforms.

download configuration cancel

```
download configuration cancel
```

Description

Cancels a scheduled incremental configuration download.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command cancels the scheduled download command completely, not just the next scheduled daily download. The `download configuration every <hour>` command must be issued again to resume automatic downloads.

If there are no downloads scheduled, this command has no effect.

Example

The following command cancels a previously scheduled download:

```
download configuration cancel
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

download configuration every

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

Description

Automatically does an incremental configuration download every day at the specified time, or immediately after switch bootup, based on the parameters specified in the `config download server` command.

Syntax Description

time	The time of day in the format <hour (0-23)>:<minutes (0-59)>.

Default

N/A.

Usage Guidelines

You must run the `config download server` command prior to using this command, to specify:

- The TFTP server and the configuration file from which the downloaded configuration will be obtained.
- Whether this TFTP server is the primary server or the secondary (backup) TFTP server.

Example

The following commands set up a scheduled incremental download of the file *config_info.txt*, to be done from the TFTP server named *tftphost* into the primary configuration area, every day at 10:00 pm:

```
config download server primary tftphost config_info.txt
download configuration every 22:00
```

History

This command was first available in ExtremeWare 6.0.

Platform Availability

This command is available on the "*i*" series platforms.

download image

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

Description

Downloads a new version of the ExtremeWare software image.

Syntax Description

ip address	Specifies the IP address of TFTP server from which the image should be obtained.
hostname	Specifies the hostname of the TFTP server from which the image should be obtained.
filename	Specifies the filename of the new image.
primary	Specifies that the new image should be stored as the primary image.
secondary	Specifies that the new image should be stored as the secondary image.

Default

Stores the downloaded image in the current location (the location used for the last reboot).

Usage Guidelines

Prior to downloading an image, you must place the new image in a file on a TFTP server on your network. Unless you include a path with the filename, this command assumes that the file resides in the same directory as the TFTP server itself.

The switch can store up to two images: a primary image and a secondary image. When you download a new image, you must select into which image space (primary or secondary) you want the new image to be placed. If no parameters are specified, the software image is saved to the current image.

Use of the `<hostname>` parameter requires that DNS be enabled.

Example

The following command downloads the switch software image from the TFTP server named *tftphost*, from the file named *s4119b2.xtr*, to the secondary image store:

```
download image tftphost s4119b2.xtr secondary
```

History

This command was available in ExtremeWare 2.0.

Support for the `<hostname>` parameter was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

save configuration

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

Description

Saves the current configuration from the switch's runtime memory to non-volatile memory.

Syntax Description

primary	Specifies the primary saved configuration.
secondary	Specifies the secondary saved configuration.

Default

Saves the current configuration to the location used on the last reboot.

Usage Guidelines

The configuration takes effect on the next reboot.

Example

The following command save the current switch configuration in the secondary configuration area:

```
save configuration secondary
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

show configuration

```
show configuration
```

Description

Displays the currently active configuration to the terminal.

Syntax Description

This command has no arguments or variables.

Usage Guidelines

If the output scrolls off the top of the screen, you can use the `enable clipaging` command to pause the display when the output fills the screen. The default for clipaging is enabled.

Example

This command shows the current configuration active in the switch:

```
show config
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

synchronize

```
synchronize
```

Description

Replicates all saved images and configurations from the master MSM to the slave MSM on the BlackDiamond.

Syntax Description

This command has no arguments or variables.

Usage Guidelines

In addition to replicating the configuration settings and images, this command also replicates which configuration or image the MSM should use on subsequent reboots.

This command does not replicate the run-time configuration. You must use the save configuration command to store the run-time configuration first. It also does not replicate the BootROM image stored on the MSM.

Example

The following command replicates all saved images and configurations from the master MSM to the slave MSM:

```
synchronize
```

History

This command was first available in ExtremeWare 4.0.

Platform Availability

This command is available on the BlackDiamond only.

unconfig switch

```
unconfig switch {all}
```

Description

Returns the switch configuration to its factory default settings.

Syntax Description

all	Specifies that the entire current configuration should be erased, and the switch rebooted.

Default

Resets configuration to factory defaults without reboot.

Usage Guidelines

Use unconfig switch to reset the configuration to factory defaults, but without erasing the configuration and rebooting. This preserves users account information, date and time settings, and so on.

Include the parameter `all` to clear the entire current configuration, including all switch parameters, and reboot using the last used image and configuration.

Example

The following command erases the entire current configuration, resets to factory defaults, and reboots the switch using the last specified saved image and saved configuration:

```
unconfig switch all
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

upload configuration

```
upload configuration [<ip address> | <hostname>] <filename> {every <time>}
```

Description

Uploads the current configuration to a TFTP server on your network.

Syntax Description

ip address	Specifies the IP address of the TFTP server.
hostname	Specifies the hostname of the TFTP server.
filename	Specifies a name for the file where the configuration is to be saved.
time	The time of day in the format <hour (0-23)>:<minutes (0-59)>.

Default

Uploads the current configuration immediately.

Usage Guidelines

The filename can be up to 255 characters long, and cannot include any spaces, commas, quotation marks, or special characters. Unless you include a path with the filename, this command places the file in the same directory as the TFTP server itself.

The uploaded ASCII file retains the command-line interface (CLI) format. This allows you to do the following:

- 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 Extreme Networks Technical Support for problem-solving purposes.

If `every <time>` is specified, the switch automatically saves the configuration to the server once per day, at the specified time. Because the filename is not changed, the configured file stored in the TFTP server is overwritten every day.

For version 4.0:

- The keyword `every` is not supported. Specify the time immediately after the filename.

For version 6.0 or later:

- The keyword `every` is required if a time is specified.

To cancel automatic upload, use the `cancel` option. If no options are specified, the current configuration is uploaded immediately.

Use of the `<hostname>` parameter requires that DNS be enabled.

Example

The following command uploads the current configuration to the file *configbackup.txt* on the TFTP server named *tftphost*, every night at 10:15 p.m.:

```
upload configuration tftphost configbackup.txt every 22:15
```

History

This command was available in ExtremeWare 2.0.

Support for the <hostname> parameter was introduced in ExtremeWare 4.0.

Platform Availability

This command is available on all platforms.

upload configuration cancel

```
upload configuration cancel
```

Description

Cancels a previously scheduled configuration upload.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

This command cancels the scheduled upload command completely, not just the next scheduled daily upload. You must re-issue the `upload configuration every <hour>` command to resume automatic uploads.

If there are no uploads scheduled, this command has no effect.

Example

The following command cancels the current automatic upload schedule:

```
upload configuration cancel
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

use configuration

```
use configuration [primary | secondary]
```

Description

Configures the switch to use a previously saved configuration on the next reboot.

Syntax Description

primary	Specifies the primary saved configuration.
secondary	Specifies the secondary saved configuration.

Default

N/A.

Usage Guidelines

The keyword “configuration” can be abbreviated to “config.”

Example

The following command specifies that the next reboot should use the primary saved configuration:

```
use configuration primary
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

use image

```
use image [primary | secondary]
```

Description

Configures the switch to use a saved image on the next reboot.

Syntax Description

primary	Specifies the primary saved software image.
secondary	Specifies the secondary saved software image.

Default

Primary.

Usage Guidelines

None.

Example

The following command configures the switch to use the primary image on the next reboot:

```
use image primary
```

History

This command was available in ExtremeWare 2.0.

Platform Availability

This command is available on all platforms.

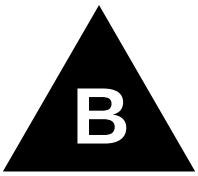

Troubleshooting Commands

If you encounter problems when using your switch, ExtremeWare provides troubleshooting commands. Use these commands only under the guidance of Extreme Networks technical personnel.

You can contact Extreme Networks technical support at (800) 998-2408 or (408) 579-2826.

If CPU utilization is high, use the debug trace commands sparingly, as they require the CPU. Disable any external syslog before you configure a debug trace, because the debug trace utility can send large amounts of information to the syslog, and if your syslog is external, that information travels over your network.

Configure a debug trace at lower levels first, and look for obvious problems. Higher levels typically record so much information that they record enough information within a few seconds.

clear debug-trace

```
clear debug-trace
```

Description

Resets the debug-trace levels to the factory settings of level 0.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

None.

Example

The following command resets the debug-trace levels to level 0:

```
clear debug-trace
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all *"i"* series platforms.

config debug-trace access-list

```
config debug-trace access-list <debug level>
```

Description

This command is not currently supported.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Not currently supported. 1 — Not currently supported. 2 — Not currently supported. 3 — Not currently supported. 4 — Not currently supported. 5 — Not currently supported.

Default

The default level is 0.

Usage Guidelines

This command is not currently supported.

Example

This command is not currently supported.

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace bgp-events

```
config debug-trace bgp-events <debug level>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — None. 1 — None. 2 — Records different stages of Finite State Machine (FSM) for BGP neighbor's negotiation process. The state of FSM transitions to next state upon events occurred. The major events in this process includes: • Connects to neighbor TCP port 179 and transitions to CONNECT state. • Passive connected and transitions to CONNECT state. • TCP connecting failed and transitions to ACTIVE state. • Connects retry on ConnectRetry timer expired, the CONN_EXP event occurs and transitions to CONNECT state. • Sends TCP_OPEN packet to TCP connected neighbor and transitions to OPENSEND state to OPENCONFIRM state and to ESTABLISHED state. • Sends NOTIFY message to peer, the STOP event occurs and transitions to IDLE state. • Receives other STOP events and transitions to IDLE state. • Receives TCP_CLOSED event and transitions to IDLE state. • Hold Timer expired for peer, the HOLD_EXP event occurs and transitions from ESTABLISHED state to IDLE state. 3 — Records information of NewState and TCP sockets. 4 — No additional information recorded. 5 — No additional information recorded.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

To limit the output to messages from a single neighbor, use the `config debug-trace bgp-neighbor` command.

Example

The following command sets the reporting level for BGP events errors to 3:

```
config debug-trace bgp events 3
```

Following is the log output at this level:

```
<DBUG:BGP> last message repeated 15 times
<DBUG:BGP> NewState ESTABLISHED
<DBUG:BGP> last message repeated 9 times
<DBUG:BGP> NewState ESTABLISHED
<INFO:SYST> serial admin: show bgp neighbor detail
<DBUG:BGP> last message repeated 7 times
<DBUG:BGP> NewState ESTABLISHED
<INFO:SYST> serial admin: show bgp neighbor
<INFO:SYST> Log cleared
<INFO:SYST> serial admin: clear log
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace bgp-keepalive

```
config debug-trace bgp-keepalive <debug level>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — None. 1 — None. 2 — Records transmitting keepalive messages to peer with peer IP address, local socket number, and the total BGP message length is printed. The keepalive message is sent every negotiated keepalive timer period. Records received messages from peer with peer IP address printed. This keepalive message is received every negotiated keepalive timer period. 3 — Records the following neighbor negotiation FSM messages while transmitting keepalive messages: • OldState — Neighbor negotiation FSM stage before keepalive message is sent. • NewState — Neighbor negotiation FSM stage after keepalive message is sent. Records the following neighbor negotiation FSM messages while receiving keepalive messages: • OldState — Neighbor negotiation FSM stage before keepalive message is received. • NewState — Neighbor negotiation FSM stage after keepalive message is received. 4 — Prints the packet in hexadecimal format when 19 bytes of the keepalive message is received. 5 — No additional information recorded.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

To limit the output to messages from a single neighbor, use the `config debug-trace bgp-neighbor` command.

Example

The following command sets the reporting level for BGP keepalive errors to 3:

```
config debug-trace bgp-keepalive 3
```

Following is the log output at this level:

```
<DBUG:BGP> NewState ESTABLISHED
<DBUG:BGP> Sending non-update to peer 10.10.0.1 socket 14 len 19
<DBUG:BGP> Queuing keepalive for peer 10.10.0.1
<DBUG:BGP> Peer 10.10.0.1 socket 14 event KEEP_EXP oldState ESTABLISHED
<DBUG:BGP> NewState ESTABLISHED
<DBUG:BGP> Sending non-update to peer 10.10.0.1 socket 14 len 19
<DBUG:BGP> Queuing keepalive for peer 10.10.0.1
<DBUG:BGP> Peer 10.10.0.1 socket 14 event KEEP_EXP oldState ESTABLISHED
<DBUG:BGP> NewState ESTABLISHED
<DBUG:BGP> Peer 10.10.0.1 socket 14 event RX_KEEP oldState ESTABLISHED
<DBUG:BGP> Rx Keepalive from peer 10.10.0.1
<DBUG:BGP> NewState ESTABLISHED
<DBUG:BGP> Sending non-update to peer 10.10.0.1 socket 14 len 19
<DBUG:BGP> Queuing keepalive for peer 10.10.0.1
<DBUG:BGP> Peer 10.10.0.1 socket 14 event KEEP_EXP oldState ESTABLISHED
<DBUG:BGP> NewState ESTABLISHED
<DBUG:BGP> Sending non-update to peer 10.10.0.1 socket 14 len 19
<DBUG:BGP> Queuing keepalive for peer 10.10.0.1
<DBUG:BGP> Peer 10.10.0.1 socket 14 event KEEP_EXP oldState ESTABLISHED
<DBUG:BGP> NewState ESTABLISHED
<DBUG:BGP> Sending non-update to peer 10.10.0.1 socket 14 len 19
<DBUG:BGP> Queuing keepalive for peer 10.10.0.1
<DBUG:BGP> Peer 10.10.0.1 socket 14 event KEEP_EXP oldState ESTABLISHED
<DBUG:BGP> NewState ESTABLISHED
<DBUG:BGP> Sending non-update to peer 10.10.0.1 socket 14 len 19
<DBUG:BGP> Queuing keepalive for peer 10.10.0.1
<DBUG:BGP> Peer 10.10.0.1 socket 14 event KEEP_EXP oldState ESTABLISHED
<INFO:SYST> serial admin: show bgp neighbor
<INFO:SYST> serial admin: configure debug-trace bgp-keepalive 3
<INFO:SYST> serial admin: configure debug-trace bgp-keepalive 2
<INFO:SYST> Log cleared
<INFO:SYST> serial admin: clear log
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace bgp-misc

```
config debug-trace bgp-misc <debug level>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — None. 1 — None. 2 — None. 3 — Checks for changes in next hops. 4 — No additional information recorded. 5 — No additional information recorded.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for BGP miscellaneous errors to 3:

```
config debug-trace bgp-misc 3
```

Following is the log output at this level:

```
02/01/2002 13:55.52 <DBUG:BGP> Checking for changes in Next Hops.
02/01/2002 13:55.46 <DBUG:BGP> last message repeated 2 times
02/01/2002 13:55.22 <DBUG:BGP> Checking for changes in Next Hops.
02/01/2002 13:55.16 <DBUG:BGP> last message repeated 2 times
02/01/2002 13:54.52 <DBUG:BGP> Checking for changes in Next Hops.
02/01/2002 13:54.44 <INFO:SYST> Log cleared
02/01/2002 13:54.44 <INFO:SYST> serial admin: clear log static
02/01/2002 13:54.44 <DBUG:BGP> Checking for changes in Next Hops.
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace bgp-msgs

```
config debug-trace bgp-msgs <debug level>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — None. 1 — Records transmitted notification message to peer with error code or error subcode printed. 2 — Records transmitted BGP message packets to peer with peer IP address, local socket number, and total BGP message length printed. The BGP message types include: • Open — sending nonupdate to peer. • Update — sending update to peer. • Notification — sending notification message to peer with error code or error subcode recorded. Transmitted keepalive messages are not recorded. Records received BGP messages from peer with peer IP address printed. The BGP message types include: • Open — Received open from peer. • Update — Received update from peer. • Notification — Received notify from peer. Received keepalive message are not recorded. 3 — Records received keepalive message, received data from peer. 4 — Prints the following packets in hexadecimal format when the messages are received: • Open. • Update. • Notification. 5 — No additional information recorded.

Default

The default level is 0.

Usage Guidelines

The debug `level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

To limit the output to messages from a single neighbor, use the `config debug-trace bgp-neighbor` command.

Example

The following command sets the reporting level for BGP messages errors to 3:

```
config debug-trace bgp-msgs 3
```

Following is the log output at this level:

```
<DBUG:BGP> last message repeated 11 times
<DBUG:BGP> Rx update from peer 10.10.0.1
<DBUG:BGP> Rcvd data from peer 10.10.0.1
<DBUG:BGP> Rx update from peer 10.10.0.1
<DBUG:BGP> Rcvd data from peer 10.10.0.1
<DBUG:BGP> Rx update from peer 10.10.0.1
<DBUG:BGP> Rcvd data from peer 10.10.0.1
<DBUG:BGP> Rx update from peer 10.10.0.1
<DBUG:BGP> Rcvd data from peer 10.10.0.1
<DBUG:BGP> last message repeated 6 times
<DBUG:BGP> Rx update from peer 10.10.0.1
<DBUG:BGP> Rcvd data from peer 10.10.0.1
<DBUG:BGP> Rx update from peer 10.10.0.1
<DBUG:BGP> Rcvd data from peer 10.10.0.1
<DBUG:BGP> Rx update from peer 10.10.0.1
<DBUG:BGP> Rcvd data from peer 10.10.0.1
<DBUG:BGP> Rx update from peer 10.10.0.1
<DBUG:BGP> Rcvd data from peer 10.10.0.1
<DBUG:BGP> Rx update from peer 10.10.0.1
<DBUG:BGP> Rcvd data from peer 10.10.0.1
<DBUG:BGP> last message repeated 10 times
<DBUG:BGP> Rx update from peer 10.10.0.1
<DBUG:BGP> Rcvd data from peer 10.10.0.1
<DBUG:BGP> Rx update from peer 10.10.0.1
<DBUG:BGP> Rcvd data from peer 10.10.0.1
<DBUG:BGP> last message repeated 2 times
<DBUG:BGP> Rx update from peer 10.10.0.1
<DBUG:BGP> Rcvd data from peer 10.10.0.1
<DBUG:BGP> Rx update from peer 10.10.0.1
<DBUG:BGP> Rcvd data from peer 10.10.0.1
<INFO:SYST> Log cleared
<INFO:SYST> serial admin: clear log
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “*i*” series platforms.

config debug-trace bgp-neighbor

```
config debug-trace bgp-neighbor <ip address> <debug level>
```

Description

This command records debug information to the syslog.

Syntax Description

ip address	Specifies the IP address of the neighbor.
debug level	Specifies a debug level. The debug level specified here limits the level of information recorded by the associated debug-trace command.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Use this command to limit the messages recorded by other BGP debug-trace commands to a single neighbor. You can use this command at any time. If you specify a neighbor, all existing BGP debug-trace configurations are limited to the neighbor you specified. If you configure an additional BGP debug-trace, that debug-trace is automatically limited to the neighbor you specified.

You can only record messages from a single neighbor; if you specify a different neighbor, only messages from that neighbor are recorded.

This command does not affect error messages recorded by the `config debug-trace bgp-misc` command, because those error messages are not related to neighbors.

To disable this command, and record messages from all neighbors, specify an IP address of 0.0.0.0.

Example

The following command limits the BGP messages recorded to only those from 10.10.10.1, and limits the levels of errors recorded to 3:

```
config debug-trace bgp-neighbor 10.10.10.1 3
```

The log output is dependent upon the BGP debug-trace configuration.

History

This command was first available in ExtremeWare 6.2.2.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace bgp-update-in

```
config debug-trace bgp-update-in <debug level>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — None. 1 — None. 2 — Records received BGP Update messages with network prefix and next-hop attribute information. Records actions triggered by the received and update message including: • Add accepted routes to BGP routing table. • Add accepted routes to global routing table. • Delete withdrawn routes from BGP routing table. • Delete withdrawn routes from routing table. Records BGP route changes due to the corresponding peer session no longer in an established state. Actions include: • Delete routes from BGP routing table. • Delete routes from global routing table. 3 — Records best route information. 4 — No additional information recorded. 5 — No additional information recorded.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

To limit the output to messages from a single neighbor, use the `config debug-trace bgp-neighbor` command.

Example

The following command sets the reporting level for BGP update in errors to 3:

```
config debug-trace bgp-update-in 3
```

Following is the log output at this level:

```
<INFO:SYST> Log cleared
<INFO:SYST> serial admin: clear log static
```

```
<DBUG:BGP> Rx nlri 192.16.196.0/255.255.254.0 from peer 10.10.0.1
<DBUG:BGP> Rx nlri 213.171.160.0/255.255.224.0 from peer 10.10.0.1
<DBUG:BGP> Rx nlri 212.35.192.0/255.255.224.0 from peer 10.10.0.1
<DBUG:BGP> Rx nlri 212.234.180.0/255.255.255.0 from peer 10.10.0.1
<DBUG:BGP> Rx nlri 195.158.9.0/255.255.255.0 from peer 10.10.0.1
<DBUG:BGP> Rx nlri 195.158.5.0/255.255.255.0 from peer 10.10.0.1
<DBUG:BGP> Rx nlri 195.158.4.0/255.255.255.0 from peer 10.10.0.1
<DBUG:BGP> Rx nlri 202.174.148.0/255.255.255.0 from peer 10.10.0.1
<DBUG:BGP> Rx nlri 195.182.224.0/255.255.224.0 from peer 10.10.0.1
<DBUG:BGP> Rx nlri 200.24.41.0/255.255.255.0 from peer 10.10.0.1
<DBUG:BGP> Rx nlri 200.24.40.0/255.255.255.0 from peer 10.10.0.1
<DBUG:BGP> Rx nlri 200.24.39.0/255.255.255.0 from peer 10.10.0.1
<DBUG:BGP> Rx nlri 200.24.38.0/255.255.255.0 from peer 10.10.0.1
<DBUG:BGP> Rx nlri 200.24.37.0/255.255.255.0 from peer 10.10.0.1
<DBUG:BGP> Rx nlri 200.24.36.0/255.255.255.0 from peer 10.10.0.1
<DBUG:BGP> Rx nlri 193.227.32.0/255.255.255.0 from peer 10.10.0.1
<DBUG:BGP> Rx nlri 193.227.27.0/255.255.255.0 from peer 10.10.0.1
<DBUG:BGP> Rx nlri 193.227.26.0/255.255.255.0 from peer 10.10.0.1
<DBUG:BGP> Rx nlri 193.227.25.0/255.255.255.0 from peer 10.10.0.1
<DBUG:BGP> Rx nlri 193.227.16.0/255.255.255.0 from peer 10.10.0.1
<DBUG:BGP> Rx nlri 193.227.15.0/255.255.255.0 from peer 10.10.0.1
<DBUG:BGP> Rx nlri 193.227.14.0/255.255.255.0 from peer 10.10.0.1
<DBUG:BGP> Rx nlri 193.227.13.0/255.255.255.0 from peer 10.10.0.1
<DBUG:BGP> Updating FDB for NLRI 193.227.15.0 255.255.255.0.
<DBUG:BGP> Add dst 193.41.197.0/24 gw 10.10.30.10 cost 6
<DBUG:BGP> Deleting dst 193.41.197.0/24 gw 10.10.30.10 cost 7
<DBUG:BGP> Updating FDB for NLRI 193.41.197.0 255.255.255.0.
<DBUG:BGP> Add dst 193.108.128.0/23 gw 10.10.30.10 cost 6
<DBUG:BGP> Deleting dst 193.108.128.0/23 gw 10.10.30.10 cost 7
<DBUG:BGP> Updating FDB for NLRI 193.108.128.0 255.255.254.0.
<DBUG:BGP> Add dst 212.41.208.0/21 gw 10.10.30.10 cost 6
<DBUG:BGP> Deleting dst 212.41.208.0/21 gw 10.10.30.10 cost 7
<DBUG:BGP> Updating FDB for NLRI 212.41.208.0 255.255.248.0.
<DBUG:BGP> Add dst 194.183.88.0/21 gw 10.10.30.10 cost 5
<DBUG:BGP> Deleting dst 194.183.88.0/21 gw 10.10.30.10 cost 6
<DBUG:BGP> Updating FDB for NLRI 194.183.88.0 255.255.248.0.
<DBUG:BGP> Add dst 213.248.0.0/20 gw 10.10.30.10 cost 7
<DBUG:BGP> Deleting dst 213.248.0.0/20 gw 10.10.30.10 cost 8
<DBUG:BGP> Updating FDB for NLRI 213.248.0.0 255.255.240.0.
<DBUG:BGP> Add dst 62.211.192.0/18 gw 10.10.30.10 cost 6
<DBUG:BGP> Deleting dst 62.211.192.0/18 gw 10.10.30.10 cost 7
<DBUG:BGP> Updating FDB for NLRI 62.211.192.0 255.255.192.0.
<DBUG:BGP> Add dst 202.28.250.0/24 gw 10.10.30.10 cost 8
<DBUG:BGP> Deleting dst 202.28.250.0/24 gw 10.10.30.10 cost 9
<DBUG:BGP> Updating FDB for NLRI 202.28.250.0 255.255.255.0.
<DBUG:BGP> Add dst 200.68.160.0/24 gw 10.10.30.10 cost 5
<DBUG:BGP> Deleting dst 200.68.160.0/24 gw 10.10.30.10 cost 8
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace bgp-update-out

```
config debug-trace bgp-update-out <debug level>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — None. 1 — None. 2 — Records transmitted BGP update messages with network prefix and peer IP address information. Also records actions after the transmission including: • Add transmitted routes to BGP neighbor advertised routing table. 3 — No additional information recorded. 4 — No additional information recorded. 5 — No additional information recorded.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

To limit the output to messages from a single neighbor, use the `config debug-trace bgp-neighbor` command.

Example

The following command sets the reporting level for BGP update out errors to 3:

```
config debug-trace bgp-update-out 3
```

Following is the log output at this level:

```
<DBUG:BGP> Processing Int change entry 203.18.38.0/255.255.255.0
<DBUG:BGP> Processing Int change entry 198.58.1.0/255.255.255.0
<DBUG:BGP> Processing Int change entry 198.58.0.0/255.255.255.0
<DBUG:BGP> Processing Int change entry 192.251.195.0/255.255.255.0
<DBUG:BGP> Processing Int change entry 193.77.24.0/255.255.255.0
<DBUG:BGP> Qing Int Chg Blk for Rt 212.46.224.0 255.255.248.0.
<DBUG:BGP> Qing Ext Chg Blk for Rt 212.46.224.0 255.255.248.0.
<DBUG:BGP> Add dst 212.46.224.0/21 gw 10.10.30.10 cost 7
<DBUG:BGP> Deleting dst 212.46.224.0/21 gw 10.10.30.10 cost 6
<DBUG:BGP> Updating FDB for NLRI 212.46.224.0 255.255.248.0.
<DBUG:BGP> Qing Int Chg Blk for Rt 216.130.64.0 255.255.224.0.
<DBUG:BGP> Qing Ext Chg Blk for Rt 216.130.64.0 255.255.224.0.
```

```
<DBUG:BGP> Add dst 216.130.64.0/19 gw 10.10.30.10 cost 6
<DBUG:BGP> Deleting dst 216.130.64.0/19 gw 10.10.30.10 cost 5
<DBUG:BGP> Updating FDB for NLRI 216.130.64.0 255.255.224.0.
<DBUG:BGP> Qing Int Chg Blk for Rt 213.242.62.0 255.255.255.0.
<DBUG:BGP> Qing Ext Chg Blk for Rt 213.242.62.0 255.255.255.0.
<DBUG:BGP> Add dst 213.242.62.0/24 gw 10.10.30.10 cost 9
<DBUG:BGP> Deleting dst 213.242.62.0/24 gw 10.10.30.10 cost 5
<DBUG:BGP> Updating FDB for NLRI 213.242.62.0 255.255.255.0.
<DBUG:BGP> Qing Int Chg Blk for Rt 195.246.44.0 255.255.255.0.
<DBUG:BGP> Qing Ext Chg Blk foge entry 212.33.128.0/255.255.224
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace bootprelay

```
config debug-trace bootprelay <debug level> vlan <vlan name>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — None. 1 — Records error messages and tracks BOOTP messages relayed. 2 — No additional information recorded. 3 — No additional information recorded. 4 — Displays a dump of each packet. 5 — No additional information recorded.
vlan name	Specifies a VLAN name.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for BOOTP relay errors to 3:

```
config debug-trace bootprelay 3
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace bridge-learning

```
config debug-trace bridge-learning <debug level>
```

Description

This command records address learning debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — None. 1 — Records warning messages, with information such as destination address, source address, and ingress port. 2 — Records informational messages, with information such as module, packet length, type of packet, flags, ingress port, and VLAN. 3 — Displays a dump of each packet. 4 — No additional information recorded. 5 — No additional information recorded.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for address learning errors to 3:

```
config debug-trace bridge-learning 3
```

Following is the log output at this level:

```
<DBUG:STP > send_bpdu: s0 port 1:8, config
<DBUG:KERN> 0x088134e474: 00 00 14 00 02 00 0f 00
<DBUG:KERN> 0x088134e464: 7d 00 00 00 00 00 80 00 00 e0 2b 8 1 7d 00 43 01
}*********+*}*C*
<DBUG:KERN> 0x088134e454: 00 26 42 42 03 00 00 00 00 00 80 0 0 00 e0 2b 81
*&BB**********+*
<DBUG:KERN> 0x088134e444: 01 80 c2 00 00 00 00 e0 2b 81 7d 0 0 81 00 e0 00
********+*}*****
<DBUG:STP > send_bpdu: s0 port 7:1, config
#
#
<DBUG:KERN> 0x088134e474: 00 00 14 00 02 00 0f 00 ********
<DBUG:KERN> 0x088134e464: 7d 00 00 00 00 00 80 00 00 e0 2b 8 1 7d 00 40 08
}*********+*}*@*
```

```
<DBUG:KERN> 0x088134e454: 00 26 42 42 03 00 00 00 00 00 80 0 0 00 e0 2b 81
*&BB**********+*
<DBUG:KERN> 0x088134e444: 01 80 c2 00 00 00 00 e0 2b 81 7d 0 0 81 00 e0 00
********+*}*****
<DBUG:STP > send_bpdu: s0 port 1:8, config
```

History

This command was first available in ExtremeWare 6.1.9.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace bridging

```
config debug-trace bridging <debug level>
```

Description

This command records layer 2 CPU processing debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — None. 1 — Records warning messages, with information such as destination address, source address, and ingress port. 2 — Records informational messages, with information such as module, packet length, type of packet, flags, ingress port, and VLAN. 3 — Displays a hex dump of each packet. 4 — No additional information recorded. 5 — No additional information recorded.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Use level 1 to detect symptoms. Use level 2 to trace the flow of a frame.

Example

The following command sets the reporting level for bridging errors to 2:

```
config debug-trace bridging 2
```

Following is the log output at this level:

```
<WARN:BRDG> bridge.c 870 PKTDROP: malformed packet:da=00:01:30:41:f9:00
sa=00:90:27:96:22:e4 inSlot=1 Ch/Subch=32(0x20)
<WARN:BRDG> bridge.c 871 Continued: inPif=1 len=%d etype=0x800 offset=0x13 inVlan=4093
<INFO:BRDG> bridge.c 795: PKTID: inPif=1 inVlan=4093 eType=0x800 offset=0x12
<INFO:BRDG> bridge.c 794: PKTID: da=00:01:30:41:f9:00  sa=00:90:27:96:22:e4 inSlot=1
Ch/Subch=32(0x20)
<INFO:BRDG> bridge.c 756: PKTIN: 1-32(0x20)/4e/800/12/4093
<INFO:BRDG> bridge.c 1532: bridgingRoutine: Got packet. inChSub: 32(0x20) flags=0x130
len=78 q=3
```

History

This command was first available in ExtremeWare 6.1.9.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace card-state-change

```
config debug-trace card-state-change <debug level>
```

Description

This command is not currently supported.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Not currently supported. 1 — Not currently supported. 2 — Not currently supported. 3 — Not currently supported. 4 — Not currently supported. 5 — Not currently supported.

Default

The default level is 0.

Usage Guidelines

This command is not currently supported.

Example

This command is not currently supported.

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace dvmrp-cache

```
config debug-trace dvmrp-cache <debug level> vlan <vlan name>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Records error messages. 1 — Records warnings. 2 — Records verbose warnings. 3 — Displays a dump of each packet. 4 — No additional information recorded. 5 — No additional information recorded.
vlan name	Specifies a VLAN name.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Use this command to trace the detailed process of adding, deleting, and modifying a multicast cache. The IP multicast cache is a hardware forwarding entry identified by a ptag index number. The following command displays the cache entries:

```
show ipmc cache [detail] <IP multicast group>
```

The trace is based on the ingress VLAN of a cache. Use this tool if the egress list of a cache is incorrect, if there are missing cache entries, or if the DVMRP task has been intermittently suspended.

Example

The following command sets the reporting level for DVMRP cache errors to 3:

```
config debug-trace dvmrp-cache 3
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace dvmrp-hello

```
config debug-trace dvmrp-hello <debug level> vlan <vlan name>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Records error messages. 1 — Records warnings. 2 — Records verbose warnings. 3 — Displays a dump of each packet. 4 — No additional information recorded. 5 — No additional information recorded.
vlan name	Specifies a VLAN name.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

This command traces all DVMRP probe messages coming into a VLAN. Use this command if switches connected to a common network have problems establishing or maintaining normal neighbor relationships.

Example

The following command sets the reporting level for DVMRP hello errors to 3:

```
config debug-trace dvmrp-hello 3
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace dvmrp-message

```
config debug-trace dvmrp-message <debug level> vlan <vlan name>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Records error messages. 1 — Records warnings. 2 — Records verbose warnings. 3 — Displays a dump of each packet. 4 — No additional information recorded. 5 — No additional information recorded.
vlan name	Specifies a VLAN name.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

This command traces the DVMRP system messages (prune, graft, and graft acknowledgement) coming into a VLAN. Use this command if a multicast stream cannot be stopped, or does not come down to the receiver after the IGMP snooping entry is verified.

Example

The following command sets the reporting level for DVMRP message errors to 3:

```
config debug-trace dvmrp-message 3
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace dvmrp-neighbor

```
config debug-trace dvmrp-neighbor <debug level> vlan <vlan name>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Records error messages. 1 — Records warnings. 2 — Records verbose warnings. 3 — Displays a dump of each packet. 4 — No additional information recorded. 5 — No additional information recorded.
vlan name	Specifies a VLAN name.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

This command traces the state of all DVMRP neighbors on a common VLAN to monitor when a neighbor is added or deleted.

Example

The following command sets the reporting level for DVMRP neighbor errors to 3:

```
config debug-trace dvmrp-neighbor 3
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace dvmrp-route

```
config debug-trace dvmrp-route <debug level> vlan <vlan name>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Records error messages. 1 — Records warnings. 2 — Records verbose warnings. 3 — Displays a dump of each packet. 4 — No additional information recorded. 5 — No additional information recorded.
vlan name	Specifies a VLAN name.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

This command records all DVMRP route report messages coming into a VLAN. Use this command if the DVMRP routing table is incorrect or unstable.

Example

The following command sets the reporting level for DVMRP route errors to 3:

```
config debug-trace dvmrp-route 3
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace dvmrp-timer

```
config debug-trace dvmrp-timer <debug level> vlan <vlan name>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Records error messages. 1 — Records warnings. 2 — Records verbose warnings. 3 — Displays a dump of each packet. 4 — No additional information recorded. 5 — No additional information recorded.
vlan name	Specifies a VLAN name.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for DVMRP timer errors to 3:

```
config debug-trace dvmrp-timer 3
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace eaps-system

```
config debug-trace eaps-system <debug level>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Records software bugs and severe errors. 1 — Records warning messages. 2 — Records changes in state, such as a failure, and changes in port status, such as a port going down. 3 — Records events that do not cause a state change and basic debug information, such as failed PDU transmission, disabled or unconfigured ports, or inactive links. 4 — Records frequently occurring events, such as timers expiring, and detailed debug information, such as sending or receiving PDUs, VLAN ID and EAPS domain of each PDU, and config values. 5 — No additional information recorded.
vlan name	Specifies a VLAN name.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for EAPS errors to 3:

```
config debug-trace eaps-system 3
```

Following is the log output at this level:

```
<INFO:SYST> serial admin: configure debug-trace eaps-system 0
<DBUG:EAPS> eaps_runtime.c 1673: Complete state unchanged, EAPS="man1"
<DBUG:EAPS> eaps_runtime.c 931: Pdu="Health-Pdu", EAPS="man1" [MAC=00:01:30:33:14:00],
RcvdSeq#=14851, CurrSeq#
<DBUG:EAPS> eaps_runtime.c 852: pdu="Health-Pdu"
<DBUG:EAPS> eaps_runtime.c 843: [DEBUG] vlanId=10, eapsdInst=0
<DBUG:EAPS> eaps.c 520: [DEBUG] Found Control Vlan. EapsInst=0
<DBUG:EAPS> eaps.c 368: [DEBUG] Wowie!! Received EAPS_PDU_MSG
<DBUG:EAPS> eaps_runtime.c 804: EAPS-PDU Transmit OK, Vlan="c1"
<DBUG:EAPS> eaps_runtime.c 779: Sending EAPS pdu out port (1:2) vlan "c1" vlanId=10
<DBUG:EAPS> eaps_runtime.c 1295: EAPS "man1" Hello Timer expired.
```

```
<DBUG:EAPS> eaps_runtime.c 1673: Complete state unchanged, EAPS="man1"
<DBUG:EAPS> eaps_runtime.c 931: Pdu="Health-Pdu", EAPS="man1" [MAC=00:01:30:33:14:00],
RcvdSeq#=14850, CurrSeq#
```

History

This command was first available in ExtremeWare 6.1.9.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace esrp-message

```
config debug-trace esrp-message <debug level> vlan <vlan name>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Displays watchdog, taskSpawn, and msgRecv errors. 1 — Displays EdpMbuf not available error, and records when sbmalloc fails on creating esrp neighbor. 2 — Records the first time the ESRP neighbor is created for each VLAN, and when the neighbor is aged out of the neighbor list. 3 — Records VLAN ESRP-PDU transmit success with time stamp, and displays check neighbor status with time stamp. 4 — Records when switch is selected as ESRP Master, and displays ESR-PDU Transmit information such as port numbers and VLAN ID. Records that the packet was received and accepted, and the neighbor was refreshed. 5 — No additional information recorded.
vlan name	Specifies a VLAN name.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for ESRP message errors to 3:

```
config debug-trace esrp-message 3
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace esrp-state-change

```
config debug-trace esrp-state-change <debug level> vlan <vlan name>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Displays state changes after 2 minutes of system restart, including cause of failover at packet receipt. Does not display state changes from neutral to slave or changes from slave to master within 2 minutes of system restart. 1 — Displays state changes from neutral to slave, and state changes from slave to master within 2 minutes of system restart. 2 — No additional information recorded. 3 — No additional information recorded. 4 — No additional information recorded. 5 — No additional information recorded.
vlan name	Specifies a VLAN name.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for ESRP state change errors to 3:

```
config debug-trace esrp-state-change 3
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace esrp-system

```
config debug-trace esrp-system <debug level>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Displays watchdog, taskSpawn, and msgRecv errors. 1 — Displays ESRP pdu Recv, bogus vlanID, and drop pkt errors, and received mismatch esrp-timer values. 2 — No additional information recorded. 3 — Displays packet receipt errors, including information such as portNum, ipaddress, vlanID, and time stamp. Also displays ESRP-aware switch information. 4 — Records power supply failures and displays packet receipt errors, including information such as group number, timer, priority, ActivePorts, and Tracked-Info. 5 — Displays a dump of each packet.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for ESRP system errors to 3:

```
config debug-trace esrp-system 3
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace fdb

```
config debug-trace fdb <debug level> vlan <vlan name>
```

Description

This command is not currently supported.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Not currently supported. 1 — Not currently supported. 2 — Not currently supported. 3 — Not currently supported. 4 — Not currently supported. 5 — Not currently supported.
vlan name	Specifies a VLAN name.

Default

The default level is 0.

Usage Guidelines

This command is not currently supported.

Example

This command is not currently supported.

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace flow-redirect

```
config debug-trace flow-redirect <debug level>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — None. 1 — Records configuration changes and unexpected code states. 2 — Records next-hop resources becoming active or inactive. 3 — No additional information recorded. 4 — No additional information recorded. 5 — No additional information recorded.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for flow redirect errors to 2:

```
config debug-trace flow-redirect 2
```

Following is the log output at this level:

```
<INFO:IPRT> redirect next hop http1 30.0.0.9 changed to up
<DBUG:SYST> i=1 Changing Nexthop fg=fffc Source=24.3.89.150 Nexthop=30.0.0.6 Nfg=fffb
<DBUG:SYST> i=0 Changing Nexthop fg=fffc Source=24.3.89.149 Nexthop=30.0.0.5 Nfg=fffa
<DBUG:SYST> i=4 Changing Nexthop fg=fffc Source=24.3.89.148 Nexthop=30.0.0.9 Nfg=ffff
<DBUG:SYST> i=3 Changing Nexthop fg=fffc Source=24.3.89.147 Nexthop=30.0.0.8 Nfg=fffe
<DBUG:SYST> i=2 Changing Nexthop fg=fffc Source=24.3.89.146 Nexthop=30.0.0.7 Nfg=fffd
<DBUG:SYST> i=1 Changing Nexthop fg=fffc Source=24.3.89.145 Nexthop=30.0.0.6 Nfg=fffb
<DBUG:SYST> i=0 Changing Nexthop fg=fffc Source=24.3.89.144 Nexthop=30.0.0.5 Nfg=fffa
<DBUG:SYST> Sag=fffc
<DBUG:SYST> Grps0 = fffa fffb fffd fffe ffff 0 0 0
<DBUG:SYST> rLBS inst=0 inUse=1 SA=24.3.89.144 sMask=fffffff 8 dPort=50
<DBUG:SYST> Looking for entries to balance in redirect 3
<DBUG:SYST> Looking for entries to balance in redirect 2
<DBUG:SYST> Looking for entries to balance in redirect 1
<DBUG:SYST> Looking for entries to balance in redirect 0
<INFO:IPRT> redirect next hop http1 30.0.0.8 changed to up <DBUG:SYST> Balancing group
ffff
```

```
<DBUG:SYST> Adding new flow for next hop ip 30.0.0.5 group fffe
<DBUG:SYST> Balancing group fffe
<DBUG:SYST> Adding new flow for next hop ip 30.0.0.5 group fffd
<DBUG:SYST> Balancing group fffd
<DBUG:SYST> Adding new flow for next hop ip 30.0.0.5 group fffb
<DBUG:SYST> Balancing group fffb
<DBUG:SYST> Looking for entries to balance in redirect 0
<DBUG:SYST> Entry Up: Adding new flow for next hop ip 30.0.0.5 group fffa
<DBUG:SYST> redirectServerListAdd 0 4
<DBUG:SYST> redirectServerListAdd 0 3
<DBUG:SYST> redirectServerListAdd 0 2
<DBUG:SYST> redirectServerListAdd 0 1
<DB UG:SYST> redirectServerListAdd 0 0
<INFO:SYST> msm-a-console admin: enable http1
<DBUG:SYST> redirectServerListDelEntry: Checking server entry 0x866c2efc 1 4
<DBUG:SYST> redirectServerListDelEntry 0x8 66c2f5c 0 4
<DBUG:SYST> redirectServerListDelEntry: Checking server entry 0x866c198c 2 4
<DBUG:SYST> redirectServerListDelEntry 0x866c19ec 0 4
<DBUG:SYST> redirectServerListDelEntry: Checking server entry 0x866c201c 3 4
<DBUG:SYST> redirectServerListDelEntry 0x866c207c 0 4
<DBUG:SYST> redirectServerListDelEntry: Freeing server entry 0x866c3efc 0 4
<DBUG:SYST> redirectServerListDelEntry 0x866c3f8c 0 4
<DBUG:SYST> Grps0 = 0 0 0 0 0 0 0 0
<DBUG:SYST> rLBS inst=0 inUse=1 SA=24.3.89.144 sMask=fffffff 8 dPort=50
<DBUG:SYST> Entry Down: Deleting sub flow for next hop ip 30.0.0.9 group fffe
<DBUG:SYST> Entry Down: Deleting sub flow for next hop ip 30.0.0.9 group fffd
<DBUG:SYST> Entry Down: Deleting sub flow for next hop ip 30.0.0.9 group fffb
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “*i*” series platforms.

config debug-trace flowstats

```
config debug-trace flowstats <debug level>
```

Description

This command records debug information to the system log.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Records error messages, such as cannot open a socket, cannot bind a socket, or cannot add or remove a flow from health-check. 1 — No additional information recorded. 2 — No additional information recorded. 3 — No additional information recorded. 4 — No additional information recorded. 5 — Displays informational messages, such as adding or deleting a flow collector. 6 — No additional information recorded. 7 — Displays debug messages such as enabling and disabling ping-check, IP address of flow collector, and port, flow collector, and flow group information for each packet, as well as a packet dump.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for flowstats errors to 3:

```
config debug-trace flowstats 3
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace health-check

```
config debug-trace health-check [<debug level> | {filter [real | virtual]
<ip address> [ftp | http | https | imap4 | ldap | nntp | pop3 | smtp |
socks | telnet | tftp | web | wildcard | www | <tcp port number>]}]
```

Description

This command records debug information to the syslog.

Syntax Description

filter	Specifies a filter.
real	Specifies a real IP address.
virtual	Specifies a virtual IP address.
ip address	Specifies the IP address.
ftp	Specifies FTP messages.
http	Specifies HTTP messages.
https	Specifies HTTPS messages.
imap4	Specifies IMAP4 messages.
ldap	Specifies LDAP messages.
nntp	Specifies NNTP messages.
pop3	Specifies POP3 messages.
smtp	Specifies SMTP messages.
socks	Specifies SOCKS messages.
telnet	Specifies Telnet messages.
tftp	Specifies TFTP messages.
web	Specifies HTTP messages.
wildcard	Specifies messages from all services.
www	Specifies HTTP messages.
debug level	Specifies a debug level. Following are the debug levels: 0 — Records unable to initialize or add a health check due to unavailable internal resources (memory, tasks, sockets, timers, or queues). 1 — Records resources becoming active or inactive, unexpected code states, and internal resources unavailable. 2 — Records resources added to or removed from health-check, configuration parameters updated, and individual health-check activity. 3 — Records more verbose health-check activity and debug messages. 4 — No additional information recorded. 5 — No additional information recorded.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Health-check debug messages apply to all resources tracked by health-check. The messages recorded are in addition to messages you have configured for other features.

You can define a filter to limit the debug messages logged. Before you define a filter, you must configure the debug level. To define a filter, you must do the following:

1. Specify a real or virtual IP address. You can specify both a real and virtual IP address in the same command line. An IP address of 0.0.0.0 will match any IP address. Messages without associated IP addresses are logged regardless of the filters you define.
2. Specify a port or service. A service of `wildcard` or a port of 0 will match any service or port number.

The filter limits the recorded messages to those concerning the IP addresses and services you specify. If you do not configure a filter, `debug-trace` records messages at the debug level you specify for every service on every IP address.

When you save your configuration, you also save your configured filter values.

Example

The following command enables level 2 debug-tracing:

```
config debug-trace health-check 2
```

The following command then configures a filter for a specific server and service:

```
config debug-trace health-check filter real 1.2.3.4 : http
```

This configuration logs health-check debug messages at levels 0, 1, and 2 for the following:

- Generic health-check messages
- ping-check for IP address 1.2.3.4
- tcp-port-check for IP 1.2.3.4 port 80 (HTTP)
- service-check for IP 1.2.3.4 port 80 (including any virtual servers that use SLB pool member 1.2.3.4 port 80)

Alternate Example

The following command enables level 2 debug-tracing:

```
config debug-trace health-check 2
```

The following command configures a filter that provides all of the information in the preceding example, and also logs service-checks specifically for the SLB virtual server (5.6.7.8 port 80) that references SLB pool member 1.2.3.4 port 80:

```
config debug-trace health-check filter real 1.2.3.4 : http virtual 5.6.7.8 : http
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace igmp-snooping

```
config debug-trace igmp-snooping <debug level>
```

Description

This command records debug information to the system log.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — None. 1 — Records warning messages, such as a memory shortage or invalid ptag. 2 — No additional information recorded. 3 — Records informational messages, such as deleting a sender, aging out an entry, or disabling IGMP snooping. 4 — Displays a dump of each IGMP and data packet. 5 — No additional information recorded.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for igmp-snooping errors to 4:

```
config debug-trace igmp-snooping 4
```

Following is the log output at this level:

```
<INFO:SYST> msm-a-console admin: configure debug-trace igmp-snooping 5 vlan all
<DBUG:IGMP> IGMPS: igmpSnoopingDataInput L2 switch data (trunk6/224.0.0.13/15.1.6.3)
<DBUG:IGMP> IGMPS: addRouterEntry trunk6: port 1:1 raddr=15.1.6.3 rp=-2039713300
<DBUG:KERN> <--- End of chain (86499000) --->
<DBUG:KERN> 0x0886499046: 00 14 00 04 3c 59 6d f3 00 00 00 00 ****<Ym*****
<DBUG:KERN> 0x0886499036: 06 03 e0 00 00 0d 20 00 35 2f 00 01 00 02 00 69     ******
*5/*****i
<DBUG:KERN> 0x0886499026: 08 00 45 c0 00 26 3d 70 00 00 01 67 86 30 0f 01
**E**&=p***g*0**
<DBUG:KERN> 0x0886499016: 01 00 5e 00 00 0d 00 01 30 0a 7e 00 81 00 e2 58
**^*****0*~****X
<DBUG:KERN> m0 @ 0x86499000:  Length=60   m_off=22   m_data=0x86499016
<DBUG:KERN> <--- Start of chain (86499000) --->
<DBUG:IGMP> IGMPS: igmpSnoopingDataInput L2 switch data (trunk4/224.0.0.13/15.1.4.2)
<DBUG:IGMP> IGMPS: addRouterEntry trunk4: port 8:1 raddr=15.1.4.2 rp=-2039722596
<DBUG:KERN> <--- End of chain (86498700) --->
```

```
<DBUG:KERN> 0x08864987a6: 00 96 fa 00 01 00 c0 a8 64 03 00 96 fa 00 ********d*****
<DBUG:KERN> 0x0886498796: 01 00 0f 01 08 03 00 96 fa 00 01 00 0f 01 06 03
****************
<DBUG:KERN> 0x0886498786: 04 02 00 96 96 00 01 00 0f 03 01 01 00 96 96 00
****************
<DBUG:KERN> 0x0886498776: 0a 00 01 00 0f 01 08 02 00 96 96 00 01 00 0f 01
****************
<DBUG:KERN> 0x0886498766: 0f 01 04 01 00 96 0a 00 01 00 0f 01 06 01 00 96
****************
<DBUG:KERN> 0x0886498756: 09 09 00 00 01 00 0f 02 01 01 00 96 0a 00 01 00
****************
<DBUG:KERN> 0x0886498746: 1e fa 01 00 0f 01 06 03 01 00 00 04 e0 00 00 00
****************
<DBUG:KERN> 0x0886498736: 04 02 e0 00 00 0d 94 04 00 00 24 00 a8 5f 75 83
**********$**_u*
<DBUG:KERN> 0x0886498726: 08 00 46 c0 00 8c e0 e2 00 00 01 67 4f 54 0f 01
**F********gOT**
<DBUG:KERN> 0x0886498716: 01 00 5e 00 00 0d 00 01 30 0a 8c 00 81 00 e1 90
**^*****0*******
<DBUG:KERN> m0 @ 0x86498700:  Length=158   m_off=22   m_data=0x86498716
<DBUG:KERN> <--- Start of chain (86498700) --->
<DBUG:IGMP> IGMPS: igmpSnoopingDataInput L2 switch data (trunk6/224.0.0.13/15.1.6.3)
<DBUG:IGMP> IGMPS: addRouterEntry trunk6: port 1:1 raddr=15.1.6.3 rp=-2039713300
<DBUG:KERN> <--- End of chain (86499400) --->
<DBUG:KERN> 0x08864994a6: 01 00 c0 a8 64 03 00 96 fa 00
****d*****
<DBUG:KERN> 0x0886499496: 08 03 00 96 fa 00 01 00 0f 01 06 03 00 96 fa 00
****************
<DBUG:KERN> 0x0886499486: 96 00 01 00 0f 03 01 01 00 96 96 00 01 00 0f 01
****************
<DBUG:KERN> 0x0886499476: 0f 01 08 02 00 96 96 00 01 00 0f 01 04 02 00 96
****************
<DBUG:KERN> 0x0886499466: 00 96 0a 00 01 00 0f 01 06 01 00 96 0a 00 01 00
****************
<DBUG:KERN> 0x0886499456: 01 00 0f 02 01 01 00 96 0a 00 01 00 0f 01 04 01
****************
<DBUG:KERN> 0x0886499446: 0f 01 06 03 01 00 00 04 e0 00 00 00 09 09 00 00
****************
<DBUG:KERN> 0x0886499436: 06 03 e0 00 00 0d 24 00 a8 5f 75 83 1e fa 01 00
******$**_u*****
<DBUG:KERN> 0x0886499426: 08 00 45 c0 00 88 3d 6e 00 00 01 67 85 d0 0f 01
**E***=n***g****
<DBUG:KERN> 0x0886499416: 01 00 5e 00 00 0d 00 01 30 0a 7e 00 81 00 e2 58
**^*****0*~****X
<DBUG:KERN> m0 @ 0x86499400:  Length=154   m_off=22   m_data=0x86499416
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace iparp

```
config debug-trace iparp <debug level> vlan <vlan name>
```

Description

This command records debug information to the syslog.

Syntax Description

debug	Specifies a debug level. Following are the debug levels: 0 — Records IP and ARP conflicts, and duplicate IP addresses. 1 — Records the following errors: • ARP interface down • No bridge for router interface • No free new entry • Filter out multicast and broadcast source address • Header too short • ARP Ethernet/IP • Invalid hw/prot length • Wrong length 2 — Records the following errors: • Router interface down • Bad IP destination • No mbuf available • Failed to ARP • SubVLAN proxy ARP disabled, replied, or ARPing • No bridge available • No ARP available • No router interface in ARPT • Loopback entry created • Suppressed re-ARP • New ARP entry for IP/MAC address • Filtering own ARP • Target matched primary, secondary, or backup 3 — No additional information recorded. 4 — No additional information recorded. 5 — No additional information recorded.
vlan name	Specifies a VLAN name.

Default

The default level is 0.

Usage Guidelines

The debug level range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for IP ARP errors to 3:

```
config debug-trace iparp 3
```

Following is the log output at this level:

```
<DBUG:SYS > arpresolve: Filled entry for (192.168.192.12,00:00:86:54:7f:2a)
<DBUG:SYS > arpresolve: START ac=0x82f3d6e0 m=0x849a6800 IP=192.168.192.12
<DBUG:SYS > arpresolve: Filled entry for (192.168.192.12,00:00:86:54:7f:2a)
<DBUG:SYS > arpresolve: START ac=0x82f3d6e0 m=0x849a6c00 IP=192.168.192.12
<DBUG:SYS > arpresolve: Filled entry for (192.168.192.12,00:00:86:54:7f:2a)
<DBUG:SYS > arpresolve: START ac=0x82f3d6e0 m=0x849a6c00 IP=192.168.192.12
<DBUG:SYS > arpresolve: Filled entry for (192.168.192.12,00:00:86:54:7f:2a)
<DBUG:SYS > arpresolve: START ac=0x82f3d6e0 m=0x849a6800 IP=192.168.192.12
<DBUG:SYS > arpresolve: Filled entry for (192.168.192.12,00:00:86:54:7f:2a)
<INFO:SYST> serial admin: configure debug-trace iparp 3 t2
<INFO:SYST> Port 2:1 link active 100Mbs FULL duplex
<INFO:SYST> serial admin: configure t2 add ports 2 : 1
<INFO:SYST> serial admin: configure t2 delete ports 1 : 1
<INFO:SYST> serial admin: enable ipforwarding t2
<INFO:SYST> serial admin: configure t2 ipaddress 192.168.192.1 / 24
<INFO:SYST> serial admin: configure t2 add ports 1 : 1
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all “*i*” series platforms.

config debug-trace ip-forwarding

```
config debug-trace ip-forwarding <debug level>
```

Description

This command records IP forwarding debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — None. 1 — Records warning messages, such as “bad checksum” or “short header length.” 2 — Records informational messages, with information such as source IP address, destination IP address, ingress port, and router interface. 3 — Displays a dump of each packet. 4 — No additional information recorded. 5 — No additional information recorded.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Use level 1 to detect symptoms. Use level 2 to trace the flow of packet processing. All levels show layer 3 information.

Example

The following command sets the reporting level for IP forwarding errors to 2:

```
config debug-trace ip-forwarding 2
```

Following is the log output at this level:

```
<WARN:IPRT> ip_slowpath.c 400: DROP: Bad IP header chksum.
source/dest=211.41.176.100/210.10.20.100
<WARN:IPRT> ip_input.c 220:DROP: Bad chksum rtif/port=2/1:1
source/dest=211.41.176.100/210.10.20.100
<WARN:IPRT> ip_slowpath.c 432:DROP: Dest is a blackhole.
source/dest=211.41.176.100/200.10.10.100
<WARN:IPRT> ip_input.c 532: DROP: Dest is a blackhole. rtif/port=2/1:1 source/dest =
211.41.176.100 / 200.10.10.100
<WARN:IPRT> ip_slowpath.c 450: DROP: Dest is a VIP and slb-type has not been
configured on the VLANs. Source/dest= 200.1.1.200 /211.41.175.1
<WARN:IPRT> ip_slowpath.c 454: DROP: Cannot find route (NETUNREACH).
source/destination=211.41.176.100/211.41.174.100
```

```
<WARN:IPRT> ip_slowpath.c 470:DROP: Dest RTIF is not forwarding.  inRtif/outRtif=2/3
source/dest=211.41.176.100/ 211.41.175.100
<WARN:IPRT> ip_input.c 560: DROP: Dest RTIF is not forwarding. inRtif/outRtif=2/3
source/dest=211.41.176.100/ 211.41.175.100
<WARN:IPRT> ip_input.c 590: DROP: Packet too small.rtif/port=2 /1:1 mlen=56
source/dest=211.41.176.100/ 211.41.175.100
<WARN:IPRT> ip_input.c 610:DROP: Bad version rtif/port=2/1:1
source/dest=211.41.176.100/211.41.175.100 ver=6
<WARN:IPRT> ip_input.c 680:DROP: Bad hdrlen rtif/port=2/1:1
source/dest=211.41.176.100/211.41.175.100 hlen=20
<WARN:IPRT> ip_input.c 710: DROP: IP too short rtif/port=2/1:1
source/dest=211.41.176.100/211.41.175.100  len/hlen46/50
<WARN:IPRT> ip_input.c 730: DROP: Bad dest IP.rtif/port=1/1:1
source/dest=211.41.176.100/0.0.0.0
<WARN:IPRT> ip_input.c 780: DROP: Bad src IP.rtif/port=3/2:1 source/dest=225.0.0.100
/211.41.176.100
<WARN:IPRT> ip_input.c 1020: DROP: Time exceeded(1).rtif/por=2/1:1
source/dest=211.41.176.100/64.10.1.100
<WARN:IPRT> ip_input.c 1120: DROP: Dest RTIF is not forwarding broadcast.
inRtif/outRtif=1/2 source/dest= 211.41.176.100/64.1.1.255
<INFO:IPRT> ip_input.c 822: Completed processing.  source/dest=20.1.1.200/10.1.1.200
<INFO:IPRT> ip_output.c 532: Send completed.  gw=10.1.1.200
<INFO:IPRT> ip_output.c 499: Calling if_output.  source/dest=20.1.1.200/10.1.1.200
gw=10.1.1.200 on rif1 (0) (0x8)
<INFO:IPRT> ip_output.c 276: Found route.  source/dest=20.1.1.200/10.1.1.200
gw=10.1.1.200
<INFO:IPRT> ip_output.c 256: Searching RT.  source/dest=20.1.1.200/10.1.1.200
<INFO:IPRT> ip_output.c 173: START: ip_output() Src=20.1.1.200 Dst=10.1.1.200
flags=0x0 forwarding=0.
<INFO:IPRT> ip_input.c 1938: Redirect suppressed.  source/dest=20.1.1.200/10.1.1.200
<INFO:IPRT> ip_input.c 1711: START: ip_forward.  source/dest=20.1.1.200/10.1.1.200
bcast=0
<INFO:IPRT> ip_input.c 802: Calling ipforward.  source/dest=20.1.1.200/10.1.1.200
<INFO:IPRT> ip_input.c 314: IPH: HL=20 Ver=4 TOS=0x0 ipLen=40 id=23036 off=16384
ttl=128 proto=0x6
<INFO:IPRT> ip_input.c 421: IPH: Dest=10.1.1.200  Source=20.1.1.200
<INFO:IPRT> ip_input.c 420: START:inRtif/inPort =1/8:9
```

History

This command was first available in ExtremeWare 6.1.9.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace ipxgns-message

```
config debug-trace ipxgns-message <debug level> vlan <vlan name>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — None. 1 — None. 2 — Verifies that IPX GNS messages are being sent and received. 3 — Verifies the contents of the messages. 4 — No additional information recorded. 5 — No additional information recorded.
vlan name	Specifies a VLAN name.

Default

The default level is 0.

Usage Guidelines

The debug level range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for IPX GNS message errors to 3:

```
config debug-trace ipxgns-message 3
```

Following is the log output at this level:

```
<DBUG:XSAP> SAP Traverse: Stuffing entry into packet
<DBUG:XSAP> SAP Traverse: Ignoring type 0278
<DBUG:XSAP> SAP Traverse: Ignoring type 026b
<DBUG:XSAP> SAP Traverse: Ignoring type 0640
<DBUG:XSAP> SAP Traverse: Ignoring type 0278
<DBUG:XSAP> SAP Traverse: Ignoring type 026b
<DBUG:XSAP> SAP Traverse: Ignoring type 0640
<DBUG:XSAP> type 0004 net: 3646f895 mac: 00:90:27:a1:44:3c socket: 1105
<DBUG:XSAP> SAP Traverse: Stuffing entry into packet
<DBUG:XSAP> last message repeated 9 times
<DBUG:XSAP> Rcv bcast GNS type(3) from (f0003606, 00:a0:c9:59:a4:5e) for service=0x4
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace ipxrip-message

```
config debug-trace ipxrip-message <debug level> vlan <vlan name>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — None. 1 — None. 2 — Verifies that ipxrip messages are being sent and received. 3 — Verifies the contents of the messages. 4 — Displays a dump of each packet. 5 — No additional information recorded.
vlan name	Specifies a VLAN name.

Default

The default level is 0.

Usage Guidelines

The debug level range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for IPX RIP message errors to 4:

```
config debug-trace ipxrip-message 4
```

Following is the log output at this level:

```
<INFO:SYST> serial admin: configure debug-trace ipxrip-message 4 ipxvlan
<INFO:SYST> Log cleared
<INFO:SYST> serial admin: clear log
<DBUG:KERN> 0x0881347d82: 00 03                                                 **
<DBUG:KERN> 0x0881347d72: 00 02 00 00 36 12 00 01 00 01 00 00 10 69 00 02
****6********i**
<DBUG:XRIP> Sending Rsp msg to f0001964:ff:ff:ff:ff:ff:ff len 18
<DBUG:XRIP> Added entry net 1069 hops 2 ticks 3 to rsp
<DBUG:XRIP> Added entry net 3612 hops 1 ticks 1 to rsp
<INFO:EAPS> eaps_runtime.c 1426: State Change, Failed -> Complete, EAPS="man1"
<INFO:EAPS> eaps_runtime.c 277: Primary Port Change, Down -> Up
<INFO:SYST> Port 1:2 link active 1000Mbs FULL duplex
<DBUG:KERN> 0x0881347d82: 00 03                                                 **
<DBUG:KERN> 0x0881347d72: 00 02 00 00 36 12 00 01 00 01 00 00 10 69 00 02
****6********i**
<DBUG:XRIP> Sending Rsp msg to f0001964:ff:ff:ff:ff:ff:ff len 18
```

```
<DBUG:XRIP> Added entry net 1069 hops 2 ticks 3 to rsp
<DBUG:XRIP> Added entry net 3612 hops 1 ticks 1 to rsp
<INFO:EAPS> eaps_runtime.c 1449: State Change, Complete -> Failed, EAPS="man1"
<INFO:EAPS> eaps_runtime.c 1018: Pdu="Link-Down-Pdu", EAPS="man1"
[MAC=00:01:30:32:ef:00]
<INFO:EAPS> eaps_runtime.c 303: Primary Port Change, Up -> Down
<INFO:SYST> Port 1:2 link down
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace ipxrip-route

```
config debug-trace ipxrip-route <debug level>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — None. 1 — None. 2 — Displays route additions and deletions. 3 — No additional information recorded. 4 — No additional information recorded. 5 — No additional information recorded.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for IPX RIP route errors to 2:

```
config debug-trace ipxrip-route 2
```

Following is the log output at this level:

```
<DBUG:XRIP> Added route to net f0220666 g/w f0001964:00:01:30:32:8d:00, hops 2, tics 2
<INFO:SYST> Log cleared
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace ipxsap-entry

```
config debug-trace ipxsap-entry <debug level>
```

Description

This command is not currently supported.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Not currently supported. 1 — Not currently supported. 2 — Not currently supported. 3 — Not currently supported. 4 — Not currently supported. 5 — Not currently supported.

Default

The default level is 0.

Usage Guidelines

This command is not currently supported.

Example

This command is not currently supported.

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace ipxsap-message

```
config debug-trace ipxsap-message <debug level> vlan <vlan name>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — None. 1 — None. 2 — Verifies that IPX SAP messages are being sent and received. 3 — Verifies the contents of the messages. 4 — Displays a dump of each packet. 5 — No additional information recorded.
vlan name	Specifies a VLAN name.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for IPX SAP message errors to 3:

```
config debug-trace ipxsap-message 3
```

Following is the log output at this level:

```
<INFO:USER> admin logged in through console
<DBUG:XSAP> Generating SAP query (opcode=0001, svc type=ffff)
<INFO:SYST> Port 2:1 link active 100Mbs FULL duplex
<INFO:SYST> Port 2:1 link down
<INFO:SYST> User admin logged out from console
<INFO:SYST> Log cleared
<INFO:SYST> serial admin: clear log
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace mpls

```
config debug-trace mpls <level>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — None. 1 — Records error and warning messages, such as session up, state machine errors, Initialization errors, label allocation errors, patricia tree failures, invalid message type or format, memory allocation errors, NVRAM parse errors, TLS tunnel creation errors, socket errors, label manager problems, and null pointer or handle. 2 — Records informational messages, such as LDP entity up, LDP parameter setting, LSP bind event, NHLFE creation, and MPLS GPP and session down errors. 3 — Records debug information, such as patricia Tree Adds/Deletes, Label Propagation and Release Msgs, Message encoding, Parameter setting, MPLS enable messages, Memory initialization, TLS setup messages, Invalid value messages, LSP Init/Teardown msgs, Event processing, RDB (route) callback information. 4 — Records more detailed debug information. 5 — No additional information recorded.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for MPLS errors to 3:

```
config debug-trace mpls 3
```

Following is the log output at this level:

```
<DBUG:IPHS> Skipping FDB refresh for 10.3.1.1 due to LSP
<DBUG:IPHS> last message repeated 3 times
<DBUG:IPHS> Skipping FDB refresh for 10.3.1.1 due to LSP
<DBUG:IPHS> last message repeated 3 times
<DBUG:IPHS> Skipping FDB refresh for 10.3.1.1 due to LSP
<DBUG:IPHS> last message repeated 3 times
<DBUG:IPHS> Skipping FDB refresh for 10.3.1.1 due to LSP
<DBUG:IPHS> last message repeated 2 times
<DBUG:IPHS> Skipping FDB refresh for 10.3.1.1 due to LSP
```

```
<DBUG:IPHS> last message repeated 2 times
<DBUG:IPHS> Skipping FDB refresh for 10.3.1.1 due to LSP
<INFO:SYST> msm-a-console admin: configure debug-trace mpls 0
<DBUG:MPLS> Slot6 MPLS: KRT CHG - Can't Add MpIdx 17 from Nh Entry 1324
<DBUG:MPLS> Slot6 MPLS: processNhlfeTableAddMpIdx: Attempting to add Mp Entry 17 to
'unused' Nhlfe Idx 1324
<DBUG:IPHS> Skipping FDB refresh for 10.3.1.1 due to LSP
<DBUG:MPLS> ip_output.c 664: CONTINUING IP OUTPUT
<INFO:MPLS> mpls_lpe.c 3277: mpls_lpe_common_input() returned
MPLS_LPE_PACKET_UNLABELLED
<DBUG:MPLS> mpls_lsp_endpt.c 346: Attempt to delete endpt entry 10.3.1.1/32
:advertise=0
<DBUG:MPLS> mpls_gpp.c 1037: MPLS Del NHLFE
<INFO:MPLS> mpls_gpp.c 1617: Create ILM for FecIp=10.3.1.1, NhlfeIx=1324,
EndptIx=1332, InLabel=0x11, OutLabel=
<DBUG:MPLS> mpls_gpp.c 796: MPLS Del ILM
<DBUG:MPLS> mpls_lsp_bind.c 1434: Unbinding label 0x00000011 from outgoing Ifc 3 Label
0x00000003
<DBUG:MPLS> mpls_rdb.c 2689: RDB REQ not able to find 10.3.1.1/32 nhop 0.0.0.0
<DBUG:MPLS> mpls_rdb.c 2524: RDB REQ - Get Recompute Next Hop
<DBUG:MPLS> mplsevnt.c 584: LDP DU LSP ID RELEASE: 0x00002356
<DBUG:MPLS> mpls_lsp_bind.c 1034: Ingress torn down for unknown LSP to endpt
10.3.1.1:32 LSP ID 0x00002356
<INFO:MPLS> mplsevnt.c 394: LMS Notify (0x8d3d1f2c): LSP TORN DOWN (5) FEC:10.3.1.1:32
nhop 0.0.0.0 LSPID:9046
<DBUG:MPLS> mpls_lsp_bind.c 466: Unbinded LSP to 10.3.1.1:32 Label 0x00000003
<DBUG:MPLS> mpls_lsp_endpt.c 765: MPLS Initiating SPF caculation for unbinded LSP to
10.3.1.1
<DBUG:MPLS> mpls_lsp_endpt.c 759: Cannot unbind LSP to 10.3.1.1:32 Type 1 nhop
10.0.1.2 without route entry
<DBUG:MPLS> mpls_rdb.c 2689: RDB REQ not able to find 10.3.1.1/32 nhop 10.0.1.2
<DBUG:MPLS> mpls_lsp_endpt.c 735: unbind_from_ipv4_endpoint: 10.3.1.1:32 Type 1 nhop
10.0.1.2
<DBUG:MPLS> mpls_rdb.c 2689: RDB REQ not able to find 10.3.1.1/32 nhop 0.0.0.0
<DBUG:MPLS> mpls_rdb.c 2524: RDB REQ - Get Recompute Next Hop
<DBUG:MPLS> mpls_rdb.c 606: mpls_rdb_callback: delete route to 17.17.17.1/32 nhop
10.0.2.2 orig=33 watch 0
<DBUG:MPLS> mpls_rdb.c 956: Recompute Issued for 10.3.1.1/32 nhop 10.0.1.2
<DBUG:MPLS> mpls_rdb.c 606: mpls_rdb_callback: delete route to 10.3.1.1/32 nhop
10.0.1.2 orig=33 watch 2
<DBUG:MPLS> mpls_rdb.c 606: mpls_rdb_callback: delete route to 192.168.100.12/32 nhop
10.0.2.2 orig=33 watch 0
<DBUG:MPLS> mpls_rdb.c 606: mpls_rdb_callback: delete route to 192.168.100.11/32 nhop
10.0.2.2 orig=33 watch 0
<DBUG:MPLS> mpls_rdb.c 606: mpls_rdb_callback: delete route to 192.168.100.2/32 nhop
10.0.2.2 orig=33 watch 0
<DBUG:MPLS> mpls_rdb.c 606: mpls_rdb_callback: delete route to 20.0.0.1/32 nhop
10.0.1.2 orig=33 watch 0
<DBUG:MPLS> mplsevnt.c 584: LDP DU LSP ID RELEASE: 0x00002371
<DBUG:MPLS> mplsevnt.c 584: LDP DU LSP ID RELEASE: 0x00002362
<DBUG:MPLS> mpls_rdb.c 606: mpls_rdb_callback: delete route to 12.0.0.1/32 nhop
10.0.2.2 orig=33 watch 0
<DBUG:MPLS> mpls_rdb.c 1060: du_recognize_new_fec: ifIndex=3,
destIp=192.168.100.11/32, nHop=10.0.2.2
<DBUG:MPLS> mpls_rdb.c 578: mpls_rdb_callback: add route to 192.168.100.11/32 nhop
10.0.2.2 watch 0 orig=33
```

```
<DBUG:MPLS> mplsevnt.c 584: LDP DU LSP ID RELEASE: 0x00002355
<DBUG:MPLS> mplsevnt.c 584: LDP DU LSP ID RELEASE: 0x00002352
<DBUG:MPLS> mplsevnt.c 584: LDP DU LSP ID RELEASE: 0x0000236D
<DBUG:MPLS> mplsevnt.c 584: LDP DU LSP ID RELEASE: 0x00002350
<DBUG:MPLS> mpls_rdb.c 1060: du_recognize_new_fec: ifIndex=3, destIp=17.17.17.1/32,
nHop=10.0.2.2
<DBUG:MPLS> mpls_rdb.c 578: mpls_rdb_callback: add route to 17.17.17.1/32 nhop
10.0.2.2 watch 0 orig=33
<DBUG:MPLS> mpls_rdb.c 1060: du_recognize_new_fec: ifIndex=3,
destIp=192.168.100.12/32, nHop=10.0.2.2
<DBUG:MPLS> mpls_rdb.c 578: mpls_rdb_callback: add route to 192.168.100.12/32 nhop
10.0.2.2 watch 0 orig=33
<INFO:MPLS> mpls_gpp.c 1617: Create ILM for FecIp=10.3.1.1, NhlfeIx=1324,
EndptIx=1332, InLabel=0x11, OutLabel=
<DBUG:MPLS> mpls_gpp.c 763: MPLS Add ILM
<DBUG:MPLS> mpls_lpe.c 1302: Bind LSP Req Label 0x00000011 to endpt 10.3.1.1:32 Type 1
<INFO:MPLS> mplsevnt.c 369: LMS Notify (0x8d3d1f2c): LSP SUCCESSFUL  FEC:10.3.1.1:32
LSPID:9046  DNS LABEL:3
<DBUG:MPLS> mpls_lsp_bind.c 1783: Binded LSP to endpt 10.3.1.1:32 nhop 10.0.1.2 Label
0x00000003
<DBUG:MPLS> mpls_gpp.c 1002: MPLS Add NHLFE
<DBUG:MPLS> mpls_lsp_nhlfe.c 366: Adding NHLFE Idx 0x0000052C to 10.3.1.1:32 nhop
10.0.1.2 label 0x00000003 to
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on ExtremeWare v6.1.8b12.

Platform Availability

This command is available on the MPLS module.

config debug-trace mpls-signalling

```
config debug-trace mpls-signalling <level>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — None. 1 — Records peer interface state msgs. 2 — No additional information recorded. 3 — Records finite state machine events 4 — No additional information recorded. 5 — Records MPLS signalling packets, such as hello packets and label mappings.

Default

The default level is 1.

Usage Guidelines

The `debug level` range is from 1 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for MPLS signalling subsystem errors to 1:

```
config debug-trace mpls-signalling 1
```

Following is the log output at this level:

```
<DBUG:IPHS> last message repeated 3 times
<DBUG:IPHS> Skipping FDB refresh for 10.3.1.1 due to LSP
<DBUG:IPHS> last message repeated 3 times
<DBUG:IPHS> Skipping FDB refresh for 10.3.1.1 due to LSP
<DBUG:IPHS> last message repeated 3 times
<DBUG:IPHS> Skipping FDB refresh for 10.3.1.1 due to LSP
<DBUG:IPHS> last message repeated 3 times
<DBUG:IPHS> Skipping FDB refresh for 10.3.1.1 due to LSP
<DBUG:IPHS> last message repeated 3 times
<DBUG:IPHS> Skipping FDB refresh for 10.3.1.1 due to LSP
<INFO:SYST> User admin logged out from telnet (100.100.105.1)
<INFO:USER> admin logged in through telnet (100.100.105.1)
<DBUG:IPHS> Skipping FDB refresh for 10.3.1.1 due to LSP
<DBUG:IPHS> Skipping FDB refresh for 10.3.1.1 due to LSP
<DBUG:IPHS> last message repeated 3 times
<DBUG:IPHS> Skipping FDB refresh for 10.3.1.1 due to LSP
<DBUG:IPHS> Skipping FDB refresh for 10.3.1.1 due to LSP
```

```
<DBUG:MSIG>  => NewSt:DU_DORMANT    OldSt:ESTABLISHED   EV:INT_DEL_UPS
<DBUG:MSIG>  => SESS: 10.0.1.1 0 Peer: 10.3.1.1 0
<DBUG:MSIG> DOWN_FSM: FEC: 10.3.1.1/32
<DBUG:MSIG>  => NewSt:ESTABLISHED   OldSt:ESTABLISHED   EV:INT_DEL_UPS
<DBUG:MSIG>  => SESS: 10.0.1.1 0 Peer: 10.3.1.1 0
<DBUG:MSIG> DOWN_FSM: FEC: 10.3.1.1/32
<DBUG:MSIG> SESS: 10.0.2.1 0 Peer: 100.100.61.1 0 NewSt:OPERATIONAL  OldSt:OPERATIONAL
EV:OTHER_MSG_RX
<DBUG:MSIG>  => NewSt:IDLE             OldSt:UPS_RLS_AWT      EV:RELEASE_REQ
<DBUG:MSIG>  => NH: 10.0.1.2 (4) SESS: 10.0.2.1 0 Peer: 100.100.61.1 0
<DBUG:MSIG> UP_FSM: FEC: 10.3.1.1/32
<DBUG:MSIG>  => NewSt:UPS_RLS_AWT      OldSt:ESTABLISHED      EV:RTE_RECOMP_REQ
<DBUG:MSIG>  => NH: 10.0.1.2 (4) SESS: 10.0.2.1 0 Peer: 100.100.61.1 0
<DBUG:MSIG> UP_FSM: FEC: 10.3.1.1/32
<DBUG:MSIG>  => NH: 10.0.1.2 (4) NewNH_NoRouteToDestination (0x2) SESS: 10.0.2.1 0
Peer: 100.100.61.1 0
<DBUG:MSIG> UP_FSM: FEC: 10.3.1.1/32
<DBUG:MSIG>  => NH: 10.0.1.2 (4) NewSt:IDLE      OldSt:EST       EV:RTE_RECOMP_REQ
<DBUG:MSIG> ING_UP_FSM: FEC: 10.3.1.1/32
<DBUG:MSIG>  => NH: 10.0.1.2 (4) NewNH_NoRouteToDestination (0x2)
<DBUG:MSIG> ING_UP_FSM: FEC: 10.3.1.1/32
<DBUG:MSIG>  => NewSt:ESTABLISHED      OldSt:RESP_AWAITED     EV:INT_DNS_MAP
<DBUG:MSIG>  => NH: 10.0.1.2 (4) SESS: 10.0.2.1 0 Peer: 100.100.61.1 0
<DBUG:MSIG> UP_FSM: FEC: 10.3.1.1/32
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on ExtremeWare v6.1.8b12.

Platform Availability

This command is available on the MPLS module.

config debug-trace ospf-hello

```
config debug-trace ospf-hello <debug level> vlan <vlan name>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — None. 1 — Displays warning message with explanations if the received hello packet is rejected by the receiving routing interface. 2 — Displays message that the routing interface is sending or receiving a hello packet. 3 — No additional information recorded. 4 — Displays Level 2 message along with the packet in hexadecimal format. 5 — No additional information recorded.
vlan name	Specifies a VLAN name.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for OSPF hello errors to 3:

```
config debug-trace ospf-hello 3
```

Following is the log output at this level:

```
<DBUG:KERN> 0x08869c5470: 03 03 03 03 ****
<DBUG:KERN> 0x08869c5460: 00 0a 02 01 00 00 00 28 0f 01 06 03 0f 01 06 01
*******(********
<DBUG:KERN> 0x08869c5450: ca 8c 00 00 00 00 00 00 00 00 00 00 ff ff ff 00
****************
<DBUG:KERN> 0x08869c5440: e0 00 00 05 02 01 00 30 01 01 01 01 00 00 00 00
*******0********
<DBUG:KERN> 0x08869c5430: 00 c0 00 44 00 00 00 00 00 59 00 00 00 00 00 00
***D*****Y******
<DBUG:OSPF> (66236)Sending Hello pkt(0x869c5430) len 68 to 224.0.0.5 if rif3
<DBUG:KERN> <--- End of chain (8649b300) --->
<DBUG:KERN> 0x088649b394: 01 01 01 01 ****
<DBUG:KERN> 0x088649b384: 00 0a 02 01 00 00 00 28 0f 01 06 03 0f 01 06 01
*******(********
```

```
12.47 <DBUG:KERN> 0x088649b374: ca 8c 00 00 00 00 00 00 00 00 00 ff ff ff 00
***************
<DBUG:KERN> 0x088649b364: e0 00 00 05 02 01 00 30 03 03 03 03 00 00 00 00
*******0********
<DBUG:KERN> 0x088649b354: 45 c0 00 30 40 4e 00 00 01 59 83 4a 0f 01 06 03
E**0@N***Y*J****
<DBUG:KERN> m0 @ 0x8649b300:  Length=68   m_off=84   m_data=0x8649b354
<DBUG:KERN> <--- Start of chain (8649b300) --->
<DBUG:OSPF> (66235) Received Hello packet from 15.1.6.3 to 224.0.0.5, len 48 proto 89
<DBUG:KERN> 0x08869c5470: 64 64 02 03 dd**
<DBUG:KERN> 0x08869c5460: 00 0a 02 01 00 00 00 28 0f 02 01 01 0f 02 01 02
*******(********
<DBUG:KERN> 0x08869c5450: 74 2a 00 00 00 00 00 00 00 00 00 ff ff ff 00
t***************
<DBUG:KERN> 0x08869c5440: e0 00 00 05 02 01 00 30 01 01 01 01 00 00 00 00
*******0********
<DBUG:KERN> 0x08869c5430: 00 c0 00 44 00 00 00 00 00 59 00 00 00 00 00 00
***D*****Y******
<DBUG:OSPF> (66233)Sending Hello pkt(0x869c5430) len 68 to 224.0.0.5 if rif5
<DBUG:KERN> <--- End of chain (86498a00) --->
<INFO:SYST> msm-a-console admin: configure debug-trace ospf-hello 0 vlan all
<INFO:SYST> msm-a-console admin: clear log static
<INFO:SYST> Log cleared
<INFO:SYST> msm-a-console admin: configure debug-trace ospf-hello 4 vlan all
<DBUG:OSPF> (66336)Sending Hello pkt(0x869c5430) len 68 to 224.0.0.5 if rif3
<DBUG:OSPF> (66335) Received Hello packet from 15.1.6.3 to 224.0.0.5, len 48 proto 89
<DBUG:OSPF> (66333)Sending Hello pkt(0x869c5430) len 68 to 224.0.0.5 if rif5
<DBUG:OSPF> (66332) Received Hello packet from 15.1.4.2 to 224.0.0.5, len 48 proto 89
<DBUG:OSPF> (66332)Sending Hello pkt(0x869c5430) len 68 to 224.0.0.5 if rif2
<DBUG:OSPF> (66330) Received Hello packet from 15.2.1.2 to 224.0.0.5, len 48 proto 89
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace ospf-lsa

```
config debug-trace ospf-lsa <debug level>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — None. 1 — None. 2 — Records type 4 and type 5 OSPF packets that are sent and received via the router. The message includes: • Delete LSAs from local Link State Database. • Add LSAs to local Link State Database. • Send unicast LS Update packets to neighbor. • Send multicast LS Update packets. • Receive unicast LS Update packets from neighbor. • Receive multicast LS Update packets. • Send LS unicast LS Acknowledgement to neighbor. • Send LS multicast LS Acknowledgement. • Receive LS unicast LS Acknowledgement from neighbor. • Receive LS multicast LS Acknowledgement. • Retransmit packets. Records OSPF LSA Batch Interval Time Stamp. 3 — Records level 2 debugging message with more detailed information. 4 — Records level 3 debugging message with the sent and received packets displayed in hexadecimal format. 5 — No additional information recorded.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for OSPF LSA errors to 3:

```
config debug-trace ospf-lsa 3
```

Following is the log output at this level:

```
<INFO:SYST> msm-a-console admin: configure debug-trace ospf-lsa 3
<INFO:SYST> Log cleared
<INFO:SYST> msm-a-console admin: clear log static
<DBUG:OSPF> Ospf LSA Batch Interval Starting at 73291.
<DBUG:OSPF> Ospf LSA Batch Interval Starting at 73261.
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace ospf-neighbor

```
config debug-trace ospf-neighbor <debug level> vlan <vlan name>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — None. 1 — None. 2 — Records any change of neighbor states. Also records actions while neighbor state transitions to another state including: • Executing DR election. • Sending and receiving database description packets. • Sending and receiving LS request packets. 3 — No additional information recorded. 4 — Displays level 2 message along with OSPF packets of type 2 and type 3 in hexadecimal format. 5 — No additional information recorded.
vlan name	Specifies a VLAN name.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for OSPF neighbor errors to 3:

```
config debug-trace ospf-neighbor 3
```

Following is the log output at this level:

```
<INFO:SYST> msm-a-console admin: configure debug-trace ospf-neighbor 3 vlan all
<INFO:SYST> Log cleared
<INFO:SYST> msm-a-console admin: clear log static
<DBUG:OSPF> electDR: circ 83349bc0 dr = 15.1.4.1 bdr = 0.0.0.0
<DBUG:OSPF> NBR change, rtid 2.2.2.2 ipa 15.1.4.2 state = INIT
<DBUG:OSPF> electDR: circ 83349bc0 dr = 15.1.4.1 bdr = 15.1.4.2
<DBUG:OSPF> NBR change, rtid 2.2.2.2 ipa 15.1.4.2 state = FULL
<DBUG:OSPF> (73522) Received LS request packet from 15.1.4.2 to 15.1.4.1, len 84 proto
89
<DBUG:OSPF> (73522)Sending LS request pkt(0x81a5f8a0) len 56 to 15.1.4.2 if rif2
```

```
<DBUG:OSPF> NBR change, rtid 2.2.2.2 ipa 15.1.4.2 state = LOADING
<DBUG:OSPF> (73522)Sending Database pkt(0x81a5f8a0) len 52 to 15.1.4.2 if rif2
<DBUG:OSPF> sending ddpkt: 02020202 seq 000126ee flags SLAVE
<DBUG:OSPF> ProcDDP nbr 2.2.2.2 seq 000126ee flags MS
<DBUG:OSPF> (73522) Received Database packet from 15.1.4.2 to 15.1.4.1, len 52 proto
89
<DBUG:OSPF> (73522)Sending Database pkt(0x81a5f8a0) len 152 to 15.1.4.2 if rif2
<DBUG:OSPF> sending ddpkt: 02020202 seq 000126ed flags SLAVE
<DBUG:OSPF> NBR change, rtid 2.2.2.2 ipa 15.1.4.2 state = EXCHANGE
<DBUG:OSPF> is slave
<DBUG:OSPF> ProcDDP nbr 2.2.2.2 seq 000126ed flags I,M,MS
<DBUG:OSPF> (73522) Received Database packet from 15.1.4.2 to 15.1.4.1, len 32 proto
89
<DBUG:OSPF> (73522)Sending Database pkt(0x81a5f8a0) len 52 to 15.1.4.2 if rif2
<DBUG:OSPF> sending ddpkt: 02020202 seq 000007f1 flags I,M,MASTER
<DBUG:OSPF> NBR change, rtid 2.2.2.2 ipa 15.1.4.2 state = EX_START
<DBUG:OSPF> electDR: circ 83349bc0 dr = 15.1.4.1 bdr = 0.0.0.0
<DBUG:OSPF> NBR change, rtid 2.2.2.2 ipa 15.1.4.2 state = INIT
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace ospf-spf

```
config debug-trace ospf-spf <debug level>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — None. 1 — None. 2 — Records SPF calculation and route table update for each area. The message includes: • Intra area SPF calculation for nonbackbone areas. • Intra area route table update for nonbackbone areas. • Intra area SPF calculation for backbone area. • Intra area route table update for backbone area. • Inter area SPF calculation for backbone area. • Inter area route table update for backbone area. • External SPF calculation. • External route table update for backbone area. 3 — Records routes that are added to and deleted from the route table as a result of SPF calculations. 4 — Records level 3 debugging message with more detailed information. 5 — No additional information recorded.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for OSPF SPF errors to 4:

```
config debug-trace ospf-spf 4
```

Following is the log output at this level:

```
<INFO:SYST> msm-a-console admin: configure debug-trace ospf-spf 3
<INFO:SYST> Log cleared
<INFO:SYST> msm-a-console admin: clear log
<DBUG:OSPF> End Extern Spf Incremental 0
<DBUG:OSPF> End Extern Route Table update area 0.0.0.0
```

```
<DBUG:OSPF> Start Extern Route Table update SPF area 0.0.0.0
<DBUG:OSPF> ExtendASB: pent=0x8388E02C, incremental=0, is_first=1
<DBUG:OSPF> Extending ASBR 3.3.3.3 is_first = 1
<DBUG:OSPF> Start Extern SPF: incremental=0, is_first=1
<DBUG:OSPF> End ospfSpfInter area 0.0.0.0 uflag 0
<DBUG:OSPF> End InterArea Route Table update area 0.0.0.0
<DBUG:OSPF> Start InterArea Route Table update SPF area 0.0.0.0
<DBUG:OSPF> ExtendABR: pent=0x8388E13C, is_first=1
<DBUG:OSPF> Extending ABR 1.1.1.1
<DBUG:OSPF> Start Inter Area SPF: Area 0.0.0.0, uflag=0, is_first=1
<DBUG:OSPF> End IntraArea Spf Area 0.0.0.0
<DBUG:OSPF> End IntraArea Route Table update area 0.0.0.0
<DBUG:OSPF> Add dst 172.16.1.0/24 gw 15.1.6.3 cost 9 area 0.0.0.0 OSPFIntra
<DBUG:OSPF> Add dst 15.2.10.0/24 gw 15.2.1.2 cost 9 area 0.0.0.0 OSPFIntra
<DBUG:OSPF> Add dst 15.2.2.0/24 gw 15.2.1.2 cost 9 area 0.0.0.0 OSPFIntra
<DBUG:OSPF> Add dst 192.168.100.0/24 gw 15.1.6.3 cost 9 area 0.0.0.0 OSPFIntra
<DBUG:OSPF> Add dst 142.168.100.0/24 gw 15.1.6.3 cost 9 area 0.0.0.0 OSPFIntra
<DBUG:OSPF> Start IntraArea Route Table update SPF area 0.0.0.0
<DBUG:OSPF> MinTent ==> pent=0x00000000
<DBUG:OSPF> ExtendTent: pent=0x833C2354
<DBUG:OSPF> AddPent: pent=0x833C2354
<DBUG:OSPF> MinTent ==> pent=0x833C2354
<DBUG:OSPF> ExtendTent: pent=0x833C25DC
<DBUG:OSPF> AddPent: pent=0x833C25DC
<DBUG:OSPF> PathFind: vid=15.2.1.1, data=255.255.255.0, ptype=0x02, tos=0
<DBUG:OSPF>  >> type=0x02 LSA={T=2, Id=15.2.1.1, AR=1.1.1.1}
<DBUG:OSPF> TryAddTent: pent=0x8388DE0C, id=15.2.1.1, data=255.255.255.0, cost=8, \
<DBUG:OSPF> Added pent 0x833c26b4 numHops 1 15.2.2.0/255.255.255.0 cost 9 cost2 0 type
3 lsai 0x0 hashbucket 16
<DBUG:OSPF>  >> LSA={lsai=<NULL>}, cost2=0
<DBUG:OSPF> AddTent: vid=15.2.2.0, data=255.255.255.0, cost=9, tmp_pent=0x8388DE0C,
type=0x03 \
<DBUG:OSPF> Adding tent 15.2.2.0/255.255.255.0 9
<DBUG:OSPF> TentFind: vid=15.2.2.0, data=255.255.255.0, ptype=0x03
<DBUG:OSPF> Start Intra Area SPF: Area 0.0.0.0, is_first=1
<DBUG:OSPF> End Extern Spf Incremental 0
<DBUG:OSPF> End Extern Route Table update area 0.0.0.0
<DBUG:OSPF> Start Extern Route Table update SPF area 0.0.0.0
<DBUG:OSPF> ExtendASB: pent=0x8388E02C, incremental=0, is_first=1
<DBUG:OSPF> Extending ASBR 3.3.3.3 is_first = 1
<DBUG:OSPF> Start Extern SPF: incremental=0, is_first=1
<DBUG:OSPF> End ospfSpfInter area 0.0.0.0 uflag 0
<DBUG:OSPF> End InterArea Route Table update area 0.0.0.0
<DBUG:OSPF> Start InterArea Route Table update SPF area 0.0.0.0
<DBUG:OSPF> ExtendABR: pent=0x8388E13C, is_first=1
<DBUG:OSPF> Extending ABR 1.1.1.1
<DBUG:OSPF> Start Inter Area SPF: Area 0.0.0.0, uflag=0, is_first=1
<DBUG:OSPF> End IntraArea Spf Area 0.0.0.0
<DBUG:OSPF> End IntraArea Route Table update area 0.0.0.0
<DBUG:OSPF> Del dst 15.3.1.0/24 gw 15.1.6.3 cost 12 area 0.0.0.0 OSPFIntra
<DBUG:OSPF> Add dst 172.16.1.0/24 gw 15.1.6.3 cost 9 area 0.0.0.0 OSPFIntra
<DBUG:OSPF> Add dst 15.2.10.0/24 gw 15.2.1.2 cost 9 area 0.0.0.0 OSPFIntra
<DBUG:OSPF> Add dst 15.2.2.0/24 gw 15.2.1.2 cost 9 area 0.0.0.0 OSPFIntra
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace pim-cache

```
config debug-trace pim-cache <debug level>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Records error messages. 1 — Records warnings. 2 — Records verbose warnings. 3 — Displays a dump of each packet. 4 — No additional information recorded. 5 — No additional information recorded.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

This command traces the detailed process of adding, deleting, and modifying a multicast cache. The IP multicast cache is a hardware forwarding entry identified by a ptag index number. The following command displays the cache entries:

```
show ipmc cache [detail] <IP multicast group>
```

The trace is based on the ingress VLAN of a cache. Use this tool if the egress list of a cache is incorrect, if there are missing cache entries, or if any multicast stream jitters.

Example

The following command sets the reporting level for PIM cache errors to 3:

```
config debug-trace pim-cache 3
```

Following is the log output at this level:

```
<INFO:SYST> msm-a-console admin: configure debug-trace pim-cache 3
<DBUG:PIM> PIM: 142.168.100.100/236.58.16.16: pimSendRegStop: dst 15.1.6.3
<DBUG:PIM> PIM: ProcRegister: sptree flow exists
<DBUG:PIM> PIM: 142.168.100.100/236.58.16.16: entry timer starting for 210
<DBUG:PIM> PIM: 142.168.100.100/236.58.16.16: fwd: extending entry's life
<DBUG:PIM> PIM: ProcRegister: NoBorder: rp:15.1.4.1 src:15.1.6.3
<DBUG:PIM> PIM: ProcRegister: 142.168.100.100/236.58.16.16
<DBUG:PIM> PIM: ProcRegister: null bit set
<DBUG:PIM> PIM: 142.168.100.100/235.49.1.6: pimSendRegStop: dst 15.1.6.3
```

```
<DBUG:PIM> PIM: ProcRegister: sptree flow exists
<DBUG:PIM> PIM: 142.168.100.100/235.49.1.6: entry timer starting for 210
<DBUG:PIM> PIM: 142.168.100.100/235.49.1.6: fwd: extending entry's life
<DBUG:PIM> PIM: ProcRegister: NoBorder: rp:15.1.6.1 src:15.1.6.3
<DBUG:PIM> PIM: ProcRegister: 142.168.100.100/235.49.1.6
<DBUG:PIM> PIM: ProcRegister: null bit set
<DBUG:PIM> PIM: 142.168.100.101/235.48.13.0: pimSendRegStop: dst 15.1.6.3
<DBUG:PIM> PIM: ProcRegister: sptree flow exists
<DBUG:PIM> PIM: 142.168.100.101/235.48.13.0: entry timer starting for 210
<DBUG:PIM> PIM: 142.168.100.101/235.48.13.0: fwd: extending entry's life
<DBUG:PIM> PIM: ProcRegister: NoBorder: rp:15.2.1.1 src:15.1.6.3
<DBUG:PIM> PIM: ProcRegister: 142.168.100.101/235.48.13.0
<DBUG:PIM> PIM: ProcRegister: null bit set
<DBUG:PIM> PIM: ProcRegister: NoBorder: rp:15.1.4.1 src:15.1.6.3
<DBUG:PIM> PIM: ProcRegister: 192.168.100.201/229.55.150.208
<DBUG:PIM> PIM: ProcRegister: NoBorder: rp:15.1.4.1 src:15.2.1.2
<DBUG:PIM> PIM: ProcRegister: 15.2.2.2/229.55.150.208
<DBUG:PIM> PIM: ProcRegister: NoBorder: rp:15.1.4.1 src:15.1.6.3
<DBUG:PIM> PIM: ProcRegister: 192.168.100.201/229.55.150.208
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace pim-hello

```
config debug-trace pim-hello <debug level> vlan <vlan name>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Records error messages. 1 — Records warnings. 2 — Records verbose warnings. 3 — Displays a dump of each packet. 4 — No additional information recorded. 5 — No additional information recorded.
vlan name	Specifies a VLAN name.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

This command traces all PIM hello messages coming into a VLAN. Use this command if switches connected to a common network have problems establishing or maintaining normal neighbor relationships.

Example

The following command sets the reporting level for PIM hello errors to 3:

```
config debug-trace pim-hello 3
```

Following is the log output at this level:

```
<INFO:SYST> msm-a-console admin: configure debug-trace pim-hello 3 vlan all
<INFO:SYST> Log cleared
<INFO:SYST> msm-a-console admin: clear log static
<DBUG:PIM> PIM: Receiving Hello pkt of len 18 from src 15.2.1.2 thro 15.2.1.1
<DBUG:PIM> PIM: Receiving Hello pkt of len 18 from src 15.1.4.2 thro 15.1.4.1
<DBUG:PIM> PIM: Xmitting Hello pkt of len 18 from src 15.2.1.1 to 224.0.0.13
<DBUG:PIM> PIM: Xmitting Hello pkt of len 18 from src 15.1.4.1 to 224.0.0.13
<DBUG:PIM> PIM: Receiving Hello pkt of len 18 from src 15.1.6.3 thro 15.1.6.1
<DBUG:PIM> PIM: Xmitting Hello pkt of len 18 from src 15.1.6.1 to 224.0.0.13
<INFO:SYST> msm-a-console admin: configure debug-trace pim-hello 3 vlan all
<INFO:SYST> Log cleared
<INFO:SYST> msm-a-console admin: clear log static
```

```
<DBUG:PIM> PIM: Xmitting Hello pkt of len 18 from src 15.1.4.1 to 224.0.0.13
<DBUG:PIM> PIM: Receiving Hello pkt of len 18 from src 15.1.6.3 thro 15.1.6.1
<DBUG:PIM> PIM: Xmitting Hello pkt of len 18 from src 15.1.6.1 to 224.0.0.13
<DBUG:PIM> PIM: Receiving Hello pkt of len 18 from src 15.2.1.2 thro 15.2.1.1
<DBUG:PIM> PIM: Receiving Hello pkt of len 18 from src 15.1.4.2 thro 15.1.4.1
<DBUG:PIM> PIM: Xmitting Hello pkt of len 18 from src 15.2.1.1 to 224.0.0.13
<DBUG:PIM> PIM: Xmitting Hello pkt of len 18 from src 15.1.4.1 to 224.0.0.13
<DBUG:PIM> PIM: Receiving Hello pkt of len 18 from src 15.1.6.3 thro 15.1.6.1
<DBUG:PIM> PIM: Xmitting Hello pkt of len 18 from src 15.1.6.1 to 224.0.0.13
<INFO:SYST> msm-a-console admin: configure debug-trace pim-hello 3 vlan all
<INFO:SYST> Log cleared
<INFO:SYST> msm-a-console admin: clear log static
<INFO:SYST> Log cleared
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace pim-message

```
config debug-trace pim-message <debug level> vlan <vlan name>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Records error messages. 1 — Records warnings. 2 — Records verbose warnings. 3 — Displays a dump of each packet. 4 — No additional information recorded. 5 — No additional information recorded.
vlan name	Specifies a VLAN name.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

This command traces all PIM system messages (join, prune, graft, graft acknowledgement, and assert for PIM-DM, and join, prune, assert, register, and register-top for PIM-SM) coming into a VLAN. Use this command if a multicast stream cannot be stopped or does not come down to the receiver after the IGMP snooping entry is verified, or if the CPU load is unexpectedly high.

Example

The following command sets the reporting level for PIM message errors to 3:

```
config debug-trace pim-message 3
```

Following is the log output at this level:

```
<INFO:SYST> msm-a-console admin: configure debug-trace pim-message 3 vlan all
<DBUG:PIM> PIM: Xmitting RP Adv(8) pkt of len 22 from src 15.2.1.1 to 15.1.6.3
<DBUG:PIM> PIM: Xmitting RP Adv(8) pkt of len 22 from src 15.1.6.1 to 15.1.6.3
<DBUG:PIM> PIM: Xmitting RP Adv(8) pkt of len 22 from src 15.1.4.1 to 15.1.6.3
<DBUG:PIM> PIM: ProcPrune: src 0.0.0.0 rp 15.1.4.1 type (*,g)
<DBUG:PIM> PIM: ProcJoinPrune: joins 0 prunes 1
<DBUG:PIM> PIM: ProcJPG: handling 235.1.1.201/255.255.255.255
<DBUG:PIM> PIM: ProcJPG: una=15.1.4.1 peerRtr=0 hold_time=210 #grp=1
<DBUG:PIM> PIM: Receiving Join/Prune(3) pkt of len 34 from src 15.1.4.2 to dst
224.0.0.13 thro 15.1.4.1
```

```
<DBUG:PIM> PIM: Receiving Bootstrap(4) pkt of len 116 from src 15.1.4.2 to dst
224.0.0.13 thro 15.1.4.1
<DBUG:PIM> PIM: Receiving Bootstrap(4) pkt of len 116 from src 15.1.6.3 to dst
224.0.0.13 thro 15.1.6.1
<DBUG:PIM> PIM: ProcJoin: src 142.168.100.101 rp 15.2.1.1 type (s,g)
<DBUG:PIM> PIM: ProcJoin: src 0.0.0.0 rp 15.2.1.1 type (*,g)
<DBUG:PIM> PIM: ProcJoinPrune: joins 2 prunes 0
<DBUG:PIM> PIM: ProcJPG: handling 235.48.13.0/255.255.255.255
<DBUG:PIM> PIM: ProcJoin: src 0.0.0.0 rp 15.2.1.1 type (*,g)
<DBUG:PIM> PIM: ProcJoinPrune: joins 1 prunes 0
<DBUG:PIM> PIM: ProcJPG: handling 224.0.1.113/255.255.255.255
<DBUG:PIM> PIM: ProcJoin: src 0.0.0.0 rp 15.1.6.1 type (*,g)
<DBUG:PIM> PIM: ProcJoinPrune: joins 1 prunes 0
<DBUG:PIM> PIM: ProcJPG: handling 227.37.32.6/255.255.255.255
<DBUG:PIM> PIM: ProcJoin: src 0.0.0.0 rp 15.1.6.1 type (*,g)
<DBUG:PIM> PIM: ProcJoinPrune: joins 1 prunes 0
<DBUG:PIM> PIM: ProcJPG: handling 227.37.32.5/255.255.255.255
<DBUG:PIM> PIM: ProcJoin: src 0.0.0.0 rp 15.1.6.1 type (*,g)
<DBUG:PIM> PIM: ProcJoinPrune: joins 1 prunes 0
<DBUG:PIM> PIM: ProcJPG: handling 227.37.32.4/255.255.255.255
<DBUG:PIM> PIM: ProcJoin: src 0.0.0.0 rp 15.2.1.1 type (*,g)
<DBUG:PIM> PIM: ProcJoinPrune: joins 1 prunes 0
<DBUG:PIM> PIM: ProcJPG: handling 227.37.32.3/255.255.255.255
<DBUG:PIM> PIM: ProcJoin: src 0.0.0.0 rp 15.2.1.1 type (*,g)
<DBUG:PIM> PIM: ProcJoinPrune: joins 1 prunes 0
<DBUG:PIM> PIM: ProcJPG: handling 227.37.32.2/255.255.255.255
<DBUG:PIM> PIM: ProcJoin: src 0.0.0.0 rp 15.2.1.1 type (*,g)
<DBUG:PIM> PIM: ProcJoinPrune: joins 1 prunes 0
<DBUG:PIM> PIM: ProcJPG: handling 227.37.32.1/255.255.255.255
<DBUG:PIM> PIM: ProcJoin: src 142.168.100.100 rp 15.1.4.1 type (s,g)
<DBUG:PIM> PIM: ProcJoin: src 142.168.100.101 rp 15.1.4.1 type (s,g)
<DBUG:PIM> PIM: ProcJPG: una=15.2.1.1 peerRtr=0 hold_time=210 #grp=12
<DBUG:PIM> PIM: Receiving Join/Prune(3) pkt of len 294 from src 15.2.1.2 to dst
224.0.0.13 thro 15.2.1.1
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace pim-neighbor

```
config debug-trace pim-neighbor <debug level> vlan <vlan name>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Records error messages. 1 — Records warnings. 2 — Records verbose warnings. 3 — Displays a dump of each packet. 4 — No additional information recorded. 5 — No additional information recorded.
vlan name	Specifies a VLAN name.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

This command traces the state of all PIM neighbors on a common VLAN to monitor if, when, or how frequently a neighbor is added or deleted.

Example

The following command sets the reporting level for PIM neighbor errors to 3:

```
config debug-trace pim-neighbor 3
```

Following is the log output at this level:

```
<INFO:SYST> msm-a-console admin: configure debug-trace pim-neighbor 3 vlan all
<INFO:SYST> Log cleared
<INFO:SYST> msm-a-console admin: clear log static
<INFO:SYST> Port 8:1 link down
<INFO:SYST> Port 8:2 link down
<INFO:SYST> Port 8:3 link down
<INFO:SYST> Port 8:4 link down
<DBUG:PIM> PIM: pimDelNbr: nbr 15.1.4.2 thro iface 15.1.4.1
<INFO:SYST> Port 8:4 link down
<INFO:SYST> Port 8:3 link down
<INFO:SYST> Port 8:2 link down
<INFO:SYST> Port 8:1 link down
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace pim-rp-mgmt

```
config debug-trace pim-rp-mgmt <debug level>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Records error messages. 1 — Records warnings. 2 — Records verbose warnings. 3 — Displays a dump of each packet. 4 — No additional information recorded. 5 — No additional information recorded.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

This command traces all RP advertisement and bootstrap messages carrying rp-set information coming into a VLAN. Use this command if RP or BSR is absent or unstable. This command is for sparse mode only.

Example

The following command sets the reporting level for PIM RP management errors to 3:

```
config debug-trace pim-rp-mgmt 3
```

Following is the log output at this level:

```
<INFO:SYST> msm-a-console admin: configure debug-trace pim-rp-mgmt 3
<INFO:SYST> Log cleared
<INFO:SYST> msm-a-console admin: clear log static
<DBUG:PIM> PIM: ProcBootstrap: Wrong iif for BSR 15.1.6.3
<DBUG:PIM> PIM: ProcBootstrap: from 15.1.4.2 in 15.1.4.1 len 56
<DBUG:PIM> PIM: rpDelEntry: 15.3.1.1 (Bootstrap)
<DBUG:PIM> PIM: ProcBootstrap: rp 15.3.1.1 no longer listed
<DBUG:PIM> PIM: rpDelEntry: 15.1.4.2 (Bootstrap)
<DBUG:PIM> PIM: ProcBootstrap: rp 15.1.4.2 no longer listed
<DBUG:PIM> PIM: rpDelEntry: 15.1.8.2 (Bootstrap)
<DBUG:PIM> PIM: ProcBootstrap: rp 15.1.8.2 no longer listed
<DBUG:PIM> PIM: rpDelEntry: 15.1.4.1 (Bootstrap)
<DBUG:PIM> PIM: ProcBootstrap: rp 15.1.4.1 no longer listed
```

```
<DBUG:PIM> PIM: rpDelEntry: 15.2.1.1 (Bootstrap)
<DBUG:PIM> PIM: ProcBootstrap: rp 15.2.1.1 no longer listed
<DBUG:PIM> PIM: rpDelEntry: 15.1.6.1 (Bootstrap)
<DBUG:PIM> PIM: ProcBootstrap: rp 15.1.6.1 no longer listed
<DBUG:PIM> PIM: rpGetEntry: 192.168.100.3
<DBUG:PIM> PIM: ProcBootstrap: rp 192.168.100.3
<DBUG:PIM> PIM: rpGetEntry: 15.1.6.3
<DBUG:PIM> PIM: ProcBootstrap: rp 15.1.6.3
<DBUG:PIM> PIM: rpGetEntry: 15.1.8.3
<DBUG:PIM> PIM: ProcBootstrap: rp 15.1.8.3
<DBUG:PIM> PIM: ProcBootstrap: grp 224.0.0.0
<DBUG:PIM> PIM: ProcBootstrap: fragment Tag 40585
<DBUG:PIM> PIM: ProcBootstrap: from 15.1.6.3 in 15.1.6.1 len 56
<DBUG:PIM> PIM: ProcBootstrap: rp 15.1.4.2 no longer listed
<DBUG:PIM> PIM: rpDelEntry: 15.1.8.2 (Bootstrap)
<DBUG:PIM> PIM: ProcBootstrap: rp 15.1.8.2 no longer listed
<DBUG:PIM> PIM: rpDelEntry: 15.2.1.1 (Bootstrap)
<DBUG:PIM> PIM: ProcBootstrap: rp 15.2.1.1 no longer listed
<DBUG:PIM> PIM: rpDelEntry: 15.1.6.1 (Bootstrap)
<DBUG:PIM> PIM: ProcBootstrap: rp 15.1.6.1 no longer listed
<DBUG:PIM> PIM: rpDelEntry: 15.1.4.1 (Bootstrap)
<DBUG:PIM> PIM: ProcBootstrap: rp 15.1.4.1 no longer listed
<DBUG:PIM> PIM: rpGetEntry: 192.168.100.3
<DBUG:PIM> PIM: ProcBootstrap: rp 192.168.100.3
<DBUG:PIM> PIM: rpGetEntry: 15.1.6.3
<DBUG:PIM> PIM: ProcBootstrap: rp 15.1.6.3
<DBUG:PIM> PIM: rpGetEntry: 15.1.8.3
<DBUG:PIM> PIM: ProcBootstrap: rp 15.1.8.3
<DBUG:PIM> PIM: ProcBootstrap: grp 224.0.0.0
<DBUG:PIM> PIM: ProcBootstrap: fragment Tag 41065
<DBUG:PIM> PIM: ProcBootstrap: from 15.1.6.3 in 15.1.6.1 len 56
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace rip-message

```
config debug-trace rip-message <debug level> vlan <vlan name>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — None. 1 — None. 2 — None. 3 — Records that the switch received a response from w.x.y.z (pier) len 24 at time.time. Records that the switch sent a response to 224.0.0.9 at time.time. 4 — Displays a dump of the RIP response. Displays a dump of the RIP response received. 5 — No additional information recorded.
vlan name	Specifies a VLAN name.

Default

The default level is 0.

Usage Guidelines

The debug `level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for RIP message errors to 3:

```
config debug-trace rip-message 3
```

Following is the log output at this level:

```
<DBUG:RIP > Sending Rsp to 224.0.0.9 at 1012569160.950000
<INFO:SYST> msm-a-console admin: configure debug-trace rip-message 3 vlan all
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace rip-route-change

```
config debug-trace rip-route-change <debug level> vlan <vlan name>
```

Description

This command is not currently supported.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Not currently supported. 1 — Not currently supported. 2 — Not currently supported. 3 — Not currently supported. 4 — Not currently supported. 5 — Not currently supported.
vlan name	Specifies a VLAN name.

Default

The default level is 0.

Usage Guidelines

This command is not currently supported.

Example

This command is not currently supported.

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace rip-triggered-update

```
config debug-trace rip-triggered-update <debug level> vlan <vlan name>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — None. 1 — None. 2 — None. 3 — Records that the switch is suppressing triggered updates for x seconds. 4 — No additional information recorded. 5 — No additional information recorded.
vlan name	Specifies a VLAN name.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for RIP triggered update errors to 3:

```
config debug-trace rip-triggered-update 3
```

Following is the log output at this level:

```
<DBUG:RIP > Suppressing triggered updates for 1 secs.
<INFO:SYST> msm-a-console admin: enable rip
<INFO:SYST> msm-a-console admin: disable rip
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace slb-3dns

```
config debug-trace slb-3dns <debug level>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Records serious errors that can cause 3DNS support to fail. This includes problems associated with system resources, invalid iQuery messages, and internal SLB and 3DNS table maintenance. 1 — Records task and or socket layer errors. These errors might indicate other more serious problems. 2 — Records informational 3DNS member change notifications, state changes, or age-outs. 3 — Decodes and displays incoming and outgoing 3DNS iQuery messages. Also displays some internal table data when the 3DNS member entries are created or updated. 4 — No additional information recorded. 5 — No additional information recorded.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for SLB 3DNS errors to 3:

```
config debug-trace slb-3dns 3
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace slb-connection

```
config debug-trace slb-connection <debug level>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Records critical failures such as insufficient memory unexpected internal state. 1 — Records unaccepted or dropped connections, as well as physical ports removed from a GoGo mode group since a failed health check and physical ports added to a GoGo mode group since a passed health check. 2 — Records GoGo mode resources that fail health check, or that change health check status from fail to pass. An associated debug level 1 message will accompany this message only if this was either the first health-check to fail on a port, or the last remaining health-check to pass on a port. 3 — Records all events associated with connecting and disconnecting resources, and some SLB configuration debugging information. 4 — No additional information recorded. 5 — No additional information recorded.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for SLB connection errors to 3:

```
config debug-trace slb-connection 3
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace slb-failover

```
config debug-trace slb-failover <debug level>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Records possible software errors such as unexpected function call failures and bad function arguments. 1 — Records configuration errors, insufficient memory, and bad data from a peer SLB switch. 2 — Records when a peer SLB switch has come up or gone down. 3 — Displays debug messages. 4 — No additional information recorded. 5 — No additional information recorded.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for SLB failover errors to 3:

```
config debug-trace slb-failover 3
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace stp-in-pdu

```
config debug-trace stp-in-pdu <debug level> mgmt <slot number>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — None. 1 — None. 2 — Records that port x:y received a config BPDU. 3 — No additional information recorded. 4 — Displays a dump of each BPDU. 5 — No additional information recorded.
Slot number	Specifies a slot number.

Default

The default level is 0.

Usage Guidelines

The debug level range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for STP in PDU errors to 4:

```
config debug-trace stp-in-pdu 4
```

Following is the log output at this level:

```
<DBUG:STP > Stpd s0 Port 3:11 Received Config Bpdu
<DBUG:STP > send_bpdu: s0 port 3:3, config
<DBUG:STP > Stpd s0 Port 3:11 Received Config Bpdu
<DBUG:STP > send_bpdu: s0 port 3:3, config
<INFO:SYST> msm-a-console admin: configure debug-trace stp-in-pdu 2
<DBUG:STP > last message repeated 3 times
<DBUG:STP > send_bpdu: s0 port 3:3, config
<DBUG:STP > last message repeated 51 times
<DBUG:STP > send_bpdu: s0 port 3:3, config
<DBUG:STP > last message repeated 51 times
<DBUG:STP > send_bpdu: s0 port 3:3, config
<INFO:SYST> Log cleared
<INFO:SYST> msm-a-console admin: clear log
<DBUG:STP > last message repeated 16 times
<DBUG:STP > Stpd s0 Port 3:11 Received Config Bpdu
<DBUG:STP > send_bpdu: s0 port 3:3, config
```

```
<DBUG:STP > Stpd s0 Port 3:11 Received Config Bpdu
<DBUG:STP > send_bpdu: s0 port 3:3, config
<DBUG:STP > Stpd s0 Port 3:11 Received Config Bpdu
<DBUG:STP > send_bpdu: s0 port 3:3, config
<DBUG:STP > Stpd s0 Port 3:11 Received Config Bpdu
<DBUG:STP > send_bpdu: s0 port 3:3, config
<DBUG:STP > Stpd s0 Port 3:11 Received Config Bpdu
<DBUG:STP > send_bpdu: s0 port 3:3, config
<DBUG:STP > Stpd s0 Port 3:11 Received Config Bpdu
<DBUG:STP > send_bpdu: s0 port 3:3, config
<DBUG:STP > Stpd s0 Port 3:11 Received Config Bpdu
<DBUG:STP > send_bpdu: s0 port 3:3, config
<DBUG:STP > Stpd s0 Port 3:11 Received Config Bpdu
<DBUG:STP > send_bpdu: s0 port 3:3, config
<DBUG:STP > Stpd s0 Port 3:11 Received Config Bpdu
<DBUG:STP > send_bpdu: s0 port 3:3, config
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace stp-out-pdu

```
config debug-trace stp-out-pdu <debug level> mgmt <slot number>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — None. 1 — None. 2 — Records that port x:y sent a config BPDU. 3 — No additional information recorded. 4 — Displays a dump of each BPDU. 5 — No additional information recorded.
Slot number	Specifies a slot number.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for STP out PDU errors to 3:

```
config debug-trace stp-out-pdu 3
```

Following is the log output at this level:

```
<DBUG:STP > Stpd s0 Port 3:11 Received Config Bpdu
<DBUG:STP > send_bpdu: s0 port 3:3, config
<DBUG:STP > Stpd s0 Port 3:11 Received Config Bpdu
<DBUG:STP > send_bpdu: s0 port 3:3, config
<DBUG:STP > Stpd s0 Port 3:11 Received Config Bpdu
<DBUG:STP > send_bpdu: s0 port 3:3, config
<DBUG:STP > Stpd s0 Port 3:11 Received Config Bpdu
<DBUG:STP > send_bpdu: s0 port 3:3, config
<DBUG:STP > Stpd s0 Port 3:11 Received Config Bpdu
<DBUG:STP > send_bpdu: s0 port 3:3, config
<DBUG:STP > Stpd s0 Port 3:11 Received Config Bpdu
<DBUG:KERN> 0x088134e474: 00 00 14 00 02 00 0f 00ssages ipxrip-message ipxrip-routeay
bridge-learning 0 00 e0 2b 81     *&BB**********+*
<DBUG:KERN> 0x088134e444: 01 80 c2 00 00 00 00 e0 2b 81 7d 0 0 81 00 e0 00
********+*}*****
<DBUG:STP > send_bpdu: s0 port 3:3, config
```

```
<DBUG:STP > Stpd s0 Port 3:11 Received Config Bpdu 02/01/2002 10:50.23 <DBUG:KERN>
0x088134e474: 00 00 14 00 02 00 0f 00 ********
<DBUG:KERN> 0x088134e464: 7d 00 00 00 00 00 80 00 00 e0 2b 81 7d 00 41 03
}*********+*}*A*
<DBUG:KERN> 0x088134e454: 00 26 42 42 03 00 00 00 00 00 80 00 00 e0 2b 81
*&BB**********+*
<DBUG:KERN> 0x088134e444: 01 80 c2 00 00 00 00 e0 2b 81 7d 00 81 00 e0 00
********+*}*****
<DBUG:STP > send_bpdu: s0 port 3:3, config
<DBUG:STP > Stpd s0 Port 3:11 Received Config Bpdu
<DBUG:KERN> 0x088134e474: 00 00 14 00 02 00 0f 00 ********
<DBUG:KERN> 0x088134e464: 7d 00 00 00 00 00 80 00 00 e0 2b 81 7d 00 41 03
}*********+*}*A*
<DBUG:KERN> 0x088134e454: 00 26 42 42 03 00 00 00 00 00 80 00 00 e0 2b 81
*&BB**********+*
<DBUG:KERN> 0x088134e444: 01 80 c2 00 00 00 00 e0 2b 81 7d 00 81 00 e0 00
********+*}*****02/01/2002 10:48.09 <DBUG:STP > send_bpdu: s0 port 3:3, config
<DBUG:STP > Stpd s0 Port 3:11 Received Config Bpdu
<DBUG:STP > send_bpdu: s0 port 3:3, config
<DBUG:STP > Stpd s0 Port 3:11 Received Config Bpdu
<DBUG:STP > send_bpdu: s0 port 3:3, config
<DBUG:STP > Stpd s0 Port 3:11 Received Config Bpdu
<DBUG:STP > send_bpdu: s0 port 3:3, config
<DBUG:STP > Stpd s0 Port 3:11 Received Config Bpdu
<DBUG:STP > send_bpdu: s0 port 3:3, config
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace udp-forwarding

```
config debug-trace udp-forwarding <debug level> vlan <vlan name>
```

Description

This command is not currently supported.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Not currently supported. 1 — Not currently supported. 2 — Not currently supported. 3 — Not currently supported. 4 — Not currently supported. 5 — Not currently supported.
vlan name	Specifies a VLAN name.

Default

The default level is 0.

Usage Guidelines

This command is not currently supported.

Example

This command is not currently supported.

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace vrrp

```
config debug-trace vrrp <debug level> vlan <vlan name>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Records critical errors, such as a task crash or an interface down. 1 — Records warning messages. 2 — Records concise packet information. 3 — Records the same information recorded in level 2, with more detail. 4 — Displays a dump of each packet. 5 — No additional information recorded.
vlan name	Specifies a VLAN name.

Default

The default level is 0.

Usage Guidelines

The `debug level` range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for VRRP to 5:

```
config debug-trace vrrp 5
```

Following is the log output at this level:

```
<DBUG:SYS > Vlan/Vrid=vlan1/1 Putting virtualMac (00:00:5e:00:01:01) into arpcom
<DBUG:SYS > Vlan/Vrid=vlan1/1 Putting systemMac (00:01:30:04:c8:00) into arpcom
```

History

This command was first available in ExtremeWare 6.2.1.

Platform Availability

This command is available on all "*i*" series platforms.

config debug-trace vrrp-hello

```
config debug-trace vrrp-hello <debug level> vlan <vlan name>
```

Description

This command records debug information to the syslog.

Syntax Description

debug level	Specifies a debug level. Following are the debug levels: 0 — Records critical errors, such as a task crash or an interface down. 1 — Records warning messages, such as incorrect address, incorrect protocol, and failed checksum. 2 — Records information such as VLAN, VRID, priority, auth-type, advert-interval, and IP address. 3 — Records the same information recorded in level 2, with more detail. 4 — Displays a dump of each packet. 5 — No additional information recorded.
vlan name	Specifies a VLAN name.

Default

The default level is 0.

Usage Guidelines

The debug level range is from 0 to 5. Higher levels record more verbose messages. Higher levels also record the messages recorded at lower levels.

Example

The following command sets the reporting level for VRRP hello messages to 4:

```
config debug-trace vrrp-hello 4
```

Following is the log output at this level:

```
<DBUG:SYS > Vlan=Default: vrrpTransmit: vrid=1,pri=255,cnt_ip_addr=1,auth_type=0
advert=1,ipaddr=10.45.208.10
<DBUG:SYS > Vlan=Default: vrrpTransmit: vrid=1,pri=255,cnt_ip_addr=1,auth_type=0
advert=1,ipaddr=10.45.208.10
<DBUG:SYS > Sending vrrp-pkt(0x8313d630) len 40 to 224.0.0.18 if rif0,
mac=00:00:5e:00:01:01
<DBUG:KERN> <--- Start of chain (84859200) --->
<DBUG:KERN> m0 @ 0x84859200:  Length=40   m_off=20   m_data=0x84859214
<DBUG:KERN> 0x0884859214: 00 00 00 28 00 00 00 00 ff 70 00 00 00 00 00 00
***(*****p******
<DBUG:KERN> 0x0884859224: e0 00 00 12 21 01 ff 01 00 01 05 c4 0a 2d d0 0a
****!********-**
<DBUG:KERN> 0x0884859234: 00 00 00 00 00 00 00 00 ********
<DBUG:KERN> <--- End of chain (84859200) --->
```

History

This command was first available in ExtremeWare 6.2.1.

Platform Availability

This command is available on all "*i*" series platforms.

config diagnostics

```
config diagnostics [extended | fastpost | normal | off]
```

Description

Runs switch diagnostics at boot-up.

Syntax Description

extended	Selects an extended diagnostic routine to run at boot-up. Takes the switch fabric and ports offline, and performs extensive ASIC, ASIC-memory, and packet loopback tests. This parameter is not supported in ExtremeWare 6.1.9 or 6.2.
fastpost	Selects fastpost diagnostic routine to run at boot-up. Takes the switch fabric offline and performs a simple ASIC test.
normal	Selects normal diagnostic routine to run at boot-up. Takes the switch fabric and ports offline, and performs a simple ASIC and packet loopback test on all the ports. This parameter is not supported in ExtremeWare 6.1.9 or 6.2.
off	Stops boot-up diagnostics.

Default

N/A.

Usage Guidelines

After you configure the boot-up diagnostics, synchronize the master and slave MSM64i using the following command:

```
synchronize
```

To run diagnostics on an I/O module, use the following command:

```
run diagnostics on [normal | extended] [<slot> | msm-a | msm-b]
```

To view results of the diagnostics test, use the following command:

```
show diagnostics
```

If the diagnostics fail, replace the module with another module of the same type.

Example

The following command configures the MSM64i to run the fastest diagnostics at boot-up:

```
config diagnostics fastpost
```

History

This command was first available in ExtremeWare 6.1.9.

Platform Availability

This command is available on "*i*" series BlackDiamond switches.

run diagnostics

```
run diagnostics [extended | normal] slot [<slot number> | msm-a | msm-b]
```

Description

Runs normal or extended diagnostics on the switch, slot, or management module.

Syntax Description

extended	Runs an extended diagnostic routine. Takes the switch fabric and ports offline, and performs extensive ASIC, ASIC-memory, and packet loopback tests.
normal	Runs a normal diagnostic routine. Takes the switch fabric and ports offline, and performs a simple ASIC and packet loopback test on all the ports.
slot number	Specifies the slot number of an I/O module. This option is available only on BlackDiamond switches.
msm-a | msm- b	Specifies the slot letter of an MSM64i. This option is available only on BlackDiamond switches.

Default

N/A.

Usage Guidelines

If you run the diagnostic routine on an I/O module, that module is taken off-line while the diagnostic test is performed. The ports on the module do not forward traffic. Once the diagnostic test is completed, the I/O module is automatically reset and becomes operational again.

NOTE

Run diagnostics when the switch can be brought off-line. The tests conducted are extensive and affect traffic that must be processed by the system CPU. The diagnostics are processed by the CPU whether you run them on an I/O or a management module.

The system watchdog timer must be disabled in order to run I/O blade diagnostics. If the watchdog timer is enabled when you enter the `run diagnostics` command, the system will display a warning and will disable the watchdog timer. When the diagnostics are complete, it will re-enable the watchdog timer. (As a side-effect, this will cause the system prompt to indicate that configuration changes have been made.)

On an I/O module, the extended diagnostic routine can require significantly more time to complete, depending on the number of ports on the module.

The normal diagnostics are short series of tests that do not test all the internal ASIC functions. On a management module, the extended diagnostic routine tests all components including the internal ASIC functions. The management module is taken off-line while the diagnostic test is performed. It is reset and operational once the test is completed.

To view results of normal or extended diagnostics tests, use the following commands:

```
show diagnostics {slot [msm-a | msm-b | <slot number>]}
```

If the results indicate that the diagnostic failed, replace the module with another module of the same type.

To configure the switch to run diagnostics on an MSM64i at boot-up, use the following command:

```
config diagnostics [extended | fastpost | normal | off]
```

Example

The following command runs extended diagnostics on the module in slot 3 of a BlackDiamond chassis:

```
run diagnostics extended slot 3
```

A warning is displayed about the impact of this test, and you have the opportunity to continue or cancel the test.

```
Running extended diagnostics will disrupt network traffic on the system.
Are you sure you want to continue?  yes/no y
```

History

This command was first available in ExtremeWare 6.1.5.

The command was modified to support the MPLS module in an ExtremeWare IP Technology Services Release based on ExtremeWare v6.1.8b12.

The command was modified to support MSM64i modules in ExtremeWare 6.1.9.

The command was modified to support Alpine and Summit switches in ExtremeWare 6.2.

Platform Availability

This command is available on "*i*" series switches.

run diagnostics packet-memory slot

```
run diagnostics packet-memory slot <slot number>
```

Description

Executes packet memory scanning for all packet memory associated with the specified I/O slot on a BlackDiamond 6808 or 6816.

Syntax Description

slot number	Specifies the slot number of an I/O module. This option is available only on BlackDiamond switches. In v 6.2.1, cannot specify an MSM.

Default

N/A.

Usage Guidelines

If you run the diagnostic routine on an I/O module, that module is taken off-line while the diagnostic test is performed. The ports on the module do not forward traffic. Once the diagnostic test is completed, the I/O module is automatically reset and becomes operational again.

NOTE

Run diagnostics when the switch can be brought off-line. The tests conducted are extensive and affect traffic that must be processed by the system CPU. The diagnostics are processed by the CPU whether you run them on an I/O or a management module.

The system watchdog timer must be disabled in order to run I/O blade diagnostics. If the watchdog timer is enabled when you enter the `run diagnostics` command, the system will display a warning and will disable the watchdog timer. When the diagnostics are complete, it will re-enable the watchdog timer. (As a side-effect, this will cause the system prompt to indicate that configuration changes have been made.)

Packet memory scanning scans the specified blade to detect single bit-related memory defects and their associated buffer locations. If packet memory defects are detected, their locations are recorded in the blade's EEPROM. Up to eight occurances can be recorded. If a defect was found during the scan process, the card is reset, the defective buffer is mapped out from further use, and the I/O card is returned to the operational state. If more than eight defects are detected, or if the defects cannot be mapped out, the card is treated as a failed card and left offline. The card should then be returned through the RMA process with Extreme Networks Technical Support.

When you enter the run diagnostic command, you are warned about any potential impacts on your switch and network (since the card will be taken offline during the diagnostic process) and you will have an opportunity to confirm or cancel the test.

In version 6.2.1, packet memory scanning and defect mapping is supported only on the G8Xi, G12SXi, F48Ti, and G8Ti modules. It is not currently supported on the MSM64i.

To show the results of a packet-memory diagnostic, use the following command:

```
show diagnostics packet-memory slot <slot number>
```

Example

The following command runs a packet-memory scan on the board in slot 4 on a BlackDiamond:

```
run diagnostics packet-memory slot 4
```

The command initially generates the following messages:

```
Running packet memory diagnostics will disrupt network traffic on card 2

>>> do this only if checksum errors are detected on the said card <<<

>>> Card can potentially be taken offline to a nonoperational state <<<

Are you sure you want to continue?  yes/no y
```

If you respond with “y” the scan proceeds.

If you run the packet-memory test on a slot that has no packet memory errors, the output from the command will be similar to the following:

```
* BD3>:17 # Starts scanning packet memory on card 4.
<diagPM-1> INFO: entering packet memory scanning for card 4
.................|.................|.................
Finished scanning packet memory for card 4 --
>>> No new defect <<<
```

If packet memory errors are detected, output similar to the following is displayed:

```
* BD3>:23 # Starts scanning packet memory on card 2.
<diagPM-1> INFO: entering packet memory scanning for card 2
.................|.................|.............
Checking Struct...has 0 entries
Received Packet
00  | 34 26 49 80 64 50 14 1f 60 54 1d a3 27 ee 5c 44
10  | 01 fd 1b 2a 15 0c 4e 79 71 c5 3c 19 1e 6b 36 83
20  | 40 39 35 79 67 2e 25 6c 7e ae 01 06 49 10 61 0e
30  | 3d da 55 9d 02 67 40 62 2a 2f 3a 64 47 dc 00 86
Transmit Packet
00  | 34 26 49 80 64 50 14 1f 60 54 1d a3 27 ee 5c 44
10  | 01 fd 1b 2a 15 0c 4e 79 71 c5 3c 19 1e 6b 36 83
20  | 44 39 35 79 67 2e 25 6c 7e ae 01 06 49 10 61 0e
30  | 3d da 55 9d 02 67 40 62 2a 2f 3a 64 47 dc 00 86

MEMID=9, recov=0, bit_position=0  , addr=101290,  entry=0...
Finished scanning packet memory for card 2 --
>>> New defect(s) detected <<<
```

History

This command was first available in ExtremeWare 6.1.5.

The command was modified to support the MPLS module in an ExtremeWare IP Technology Services Release based on ExtremeWare v6.1.8b12.

The command was modified to support MSM64i modules in ExtremeWare 6.1.9.

The command was modified to support Alpine and Summit switches in ExtremeWare 6.2.

The command was modified to support Packet Memory scanning for Black Diamond I/O blades in ExtremeWare 6.2.1. See the Release Note for information on which blades are supported.

Platform Availability

This command is available on "*i*" series switches.

show debug-trace

```
show debug-trace [access-list | bgp-events | bgp-keepalive | bgp-misc |
bgp-msgs | bgp-neighbor | bgp-update-in | bgp-update-out | bootprelay |
card-state-change | dvmrp-cache | dvmrp-hello | dvmrp-message |
dvmrp-neighbor | dvmrp-route | dvmrp-timer | esrp-message |
esrp-state-change | esrp-system | fdb | flow-redirect | health-check |
iparp | ipxrip-message | ipxrip-route | ipxgns-messages | ipxsap-message |
ipxsap-entry | ospf-hello | ospf-lsa | ospf-neighbor | ospf-spf | pim-cache
| pim-hello | pim-neighbor | pim-message | pim-rp-mgmt | rip-message |
rip-route-change | rip-triggered-update | slb-connection | slb-failover |
slb-3dns | stp-in-pdu | stp-out-pdu | udp-forwarding] vlan <vlan name>
```

Description

Displays the configured debug-trace levels.

Syntax Description

access-list	Not currently supported.
bgp-events	Specifies BGP events level.
bgp-keepalive	Specifies BGP keepalive level.
bgp-misc	Specifies miscellaneous BGP level.
bgp-msgs	Specifies BGP message level.
bgp-neighbor	Not currently supported.
bgp-update-in	Specifies incoming BGP update level.
bgp-update-out	Specifies outgoing BGP update level.
bootprelay	Specifies BOOTP relay level.
card-state-change	Not currently supported.
dvmrp-cache	Specifies DVMRP cache level.
dvmrp-hello	Specifies DVMRP hello level.
dvmrp-message	Specifies DVMRP message level.
dvmrp-neighbor	Specifies DVMRP neighbor level.
dvmrp-route	Specifies DVMRP route level.
dvmrp-timer	Specifies DVMRP timer level.
esrp-message	Specifies ESRP message level.
esrp-state-change	Specifies ESRP state-change level.
esrp-system	Specifies ESRP system level.
fdb	Not currently supported.
flow-redirect	Specifies flow redirect level.
health-check	Specifies health check level.
iparp	Specifies IP ARP level.
ipxrip-message	Specifies IPX RIP message level.
ipxrip-route	Specifies IPX RIP route level.
ipxgns-messages	Specifies IPX GNS message level.
ipxsap-message	Specifies IPX SAP message level.

ipxsap-entry	Not currently supported.
ospf-hello	Specifies OSPF hello level.
ospf-lsa	Specifies OSPF LSA level.
ospf-neighbor	Specifies OSPF neighbor level.
ospf-spf	Specifies OSPF SPF level.
pim-cache	Specifies PIM cache level.
pim-hello	Specifies PIM hello level.
pim-neighbor	Specifies PIM neighbor level.
pim-message	Specifies PIM message level.
pim-rp-mgmt	Specifies PIM RP level.
rip-message	Specifies RIP message level.
rip-route-change	Specifies RIP route level.
rip-triggered-update	Specifies RIP triggered update level.
slb-connection	Specifies SLB connection level.
slb-failover	Specifies SLB failover level.
slb-3dns	Specifies SLB 3DNS level.
stp-in-pdu	Specifies incoming STP PDU level.
stp-out-pdu	Specifies outgoing STP PDU level.
udp-forwarding	Not currently supported.
vlan name	Specifies a VLAN name.

Example

The following command displays the debug-trace level configured for esrp messages:

```
show debug-trace esrp-message
```

History

This command was first available in ExtremeWare 6.1.

Platform Availability

This command is available on all "*i*" series platforms.

show diagnostics backplane mpls mapping

```
show diagnostics backplane mpls mapping {active}
```

Description

Displays diagnostic information related to the MPLS module internal backplane switch ports. This command also displays the external I/O port to internal MPLS module backplane switch port mappings.

Syntax Description

active	Specifies to limit the port mapping display to active external I/O ports only.

Default

N/A.

Usage Guidelines

This command is only supported when the backplane load-sharing policy mode is port-based. If the active parameter is specified, the port mapping display is limited to active external I/O ports only. Used in conjunction with the `show diagnostics backplane utilization` command, these commands are helpful for diagnosing over-subscription problems related to backplane I/O port switch mappings.

Example

The following command displays diagnostic information related to the MPLS module internal backplane switch ports:

```
show diagnostics backplane mpls mapping
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on ExtremeWare v6.1.8b12.

Platform Availability

This command is available on the MPLS module in the BlackDiamond switch.

show diagnostics backplane utilization

```
show diagnostics backplane utilization
```

Description

Displays backplane link utilization information.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Displays information including:

- Real-time traffic utilization on configured backplane links between active modules and MSM64i modules.
- The number of packets transmitted and received.
- The percentage of bandwidth used on the link.

Backplane utilization statistics can be reset by pressing 0 while the information is being displayed.

Example

The following command displays backplane link utilization information:

```
show diagnostics backplane utilization
```

History

This command was first available in an ExtremeWare IP Technology Services Release based on ExtremeWare v6.1.8b12.

Platform Availability

This command is available on the MPLS module in the BlackDiamond switch.

show diagnostics packet-memory slot

```
show diagnostics packet-memory slot <slot number>
```

Description

Displays the results of the packet memory scan on BlackDiamond 6808 and BlackDiamond 6816 I/O modules.

Syntax Description

slot number	Specifies the slot number of an I/O module.

Default

N/A.

Usage Guidelines

Use this command to display the results of a packet memory scan. The command output displays the number of defects identified, and the number that were recoverable. If packet memory defects were found, it displays information about each defect.

In Version 6.2.1, this applies only to the G8Xi, G8Ti, G12SXi, and F48Ti. MSM blades are not supported in this release.

Example

The following command displays the results of a PM scan for slot 2, where a single defect was found:

```
show diagnostics packet-memory slot 2
```

If no defects are found, the output will look similar to the following:

```
-----------------------------------------------------
         Packet memory defect info for card 1
-----------------------------------------------------

Num of defects = 0, num of recoverable defects = 0


-----------------------------------------------------
```

If defects are found, the output displays the number of defects, and provides information about each identified defect.

```
-----------------------------------------------------
         Packet memory defect info for card 2
-----------------------------------------------------

Num of defects = 1, num of recoverable defects = 1

Defect information:

Defect entry 1
```

```
        fault_2 = 0
        inhibit = 0
        recoverable = 0
        mem ID = 9
        bit position = 0
        address = 0x18baa

-------------------------------------------------------
```

History

This command was first available in ExtremeWare 6.2.1.

Platform Availability

This command is available on "*i*" series BlackDiamond switches only.

show diagnostics

```
show diagnostics {slot [<slot number> | msm-a | msm-b]}
```

Description

Displays the status of the system health checker as well as information from the last diagnostic test run on the switch.

Syntax Description

slot number	Specifies the slot number of an I/O module. (6.1 or higher)
msm-a | msm- b	Specifies the MSM64i. (6.1 or higher)

Default

N / A.

Usage Guidelines

Use this command to display the status of the system health checker as well as information from the last diagnostic test run on the switch. The switch diagnostics are displayed in a tabular format with the day, month, date, year, and time of the diagnostic test at the top of the table.

Table 24: Show Diagnostics Command Field Definitions

Field	Definitions
System Platform	Specifies system type (4.x only).
System Part No.	Specifies system part number, revision level, and serial number.(4.x only)
Main Board No.	Specifies main board part number, revision level, and serial number. (4.x only)
MAC Address	Specifies system MAC address. (4.x only)
Slot	Specifies the slot for which the results are displayed.
CPU System	Indicates diagnostic results.
Registers Test	Indicates diagnostic results.
Memory Test	Indicates diagnostic results.
System Test	Indicates diagnostic results.

To run diagnostics on a I/O module or MSM64i, use the following command:

```
run diagnostics [extended | normal] slot [msm-a | msm-b | <slot number>]
```

Depending on the software version running on your switch or the model of your switch, additional or different diagnostics information might be displayed.

Example

The following command displays the results of module diagnostics for slot *msm-b*:

```
show diagnostics slot msm-b
```

The results are similar to the following:

```
--------------------------------------------------------------------------
        Diagnostic Test Result run on Thu Jan 31 14:59:26 2002


--------------------------------------------------------------------------
Slot           :   B
--------------------------------------------------------------------------
CPU System     |   Passed
--------------------------------------------------------------------------
Registers Test |   Passed
--------------------------------------------------------------------------
Memory Test    |   Passed
--------------------------------------------------------------------------
System Test    |   Passed
--------------------------------------------------------------------------
```

The following command shows the results of diagnostics run on a stand-alone "*i*" series switch:

```
show diagnostics
```

The results are similar to the following:

```
--------------------------------------------------------------------------
        Diagnostic Test Result run on Thu Sep 14 16:01:15 2000


--------------------------------------------------------------------------
CPU System     |   Passed
--------------------------------------------------------------------------
Registers Test |   Passed
--------------------------------------------------------------------------
Memory Test    |   Passed
--------------------------------------------------------------------------
System Test    |   Passed
--------------------------------------------------------------------------
```

History

This command was available in ExtremeWare 4.1.19, and in ExtremeWare 6.1.5.

This command was modified in an ExtremeWare IP Technology Services Release to support MPLS and PoS modules.

The command was modified to include MSM64i modules in ExtremeWare 6.1.9.

The command was modified to support Alpine and Summit switches in ExtremeWare 6.2.

Platform Availability

This command is available on all platforms.

show tech-support

```
show tech-support
```

Description

Displays the output of various show commands to assist in monitoring and troubleshooting the switch.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

The `show tech-support` command displays the output for the following show commands:

- show version
- show switch
- show config
- show diag
- show slot
- show fdb
- show iparp
- show ipfdb
- show ipstats
- show iproute
- show ipmc cache detail
- show ipmc fdb
- show igmp snooping detail
- show memory detail
- show log

It also displays the output from internal debug commands. This command disables the CLI paging feature.

This information can be useful for your technical support representative if you experience a problem.

Depending on the software version running on your switch, additional or different show command output is displayed.

Example

The following command displays the show command output on the switch:

```
show tech-support
```

History

This command was first available in ExtremeWare 6.1.9.

Platform Availability

This command is available on all "*i*" series platforms.

top

```
top
```

Description

Displays real-time CPU utilization information by process.

Syntax Description

This command has no arguments or variables.

Default

N/A.

Usage Guidelines

Use this command to show the percentage of CPU processing devoted to each task, sampled every 30 seconds. In a healthy ExtremeWare system, only the BGTask takes up significant CPU processing power. Investigate tasks showing consistent or periodic high CPU utilization.

You can change the display by typing a character while the display is active. These single character commands are as follows:

Table 25: TOP command display options

u	Go up one screen
d	Go down one screen
c	Clear max utilization
%	Sort tasks by CPU utilization
t	Sort tasks by task id
p	Sort tasks by program counter
n	Sort tasks by name
s	Sort tasks by task state
m	Sort tasks by max CPU utilization
h	Show the help screen
<space>	go to next sort type
q <esc> <return>	Exit top

The following table defines the tasks. Depending on your switch model and the functions it is executing, you will see only a subset of these tasks.

Table 26: ExtremeWare Task Descriptions

Task	Description
httpd	The HTTP daemon task manages the HTTP web management interface on the system.
Logpoll	In an active dual CPU system, the master CPU will initiate the log polling task (Logpoll) to periodically poll the secondary or slave CPU(s). This process clears the individual syslogs and consolidates them onto the master CPU switch log.
mportTask	The management port task.
pifstate	The port interface state task (pifstate) processes port link state changes. It is watchdog timer poll driven as opposed to interrupt driven by hardware events.
tAsyncSave	The tAsyncSave tasks the NVRAM asynchronous save/write processing task. This process manages the save or writes to the NVRAM.
tbgpTask	The border gateway protocol task (tbgpTask) implements and processes BGP on the switch.
tbgpTimerTask	The BGP internal process timer task (tbgpTimerTask) manages the internal BGP timer delays for checking BGP networks and next hops.
tBgQosMon	The background Quality of Service monitor task (tBgQosMon) is a background version of the QoS monitoring task that monitors transmit count and kill count of ports and cycles as long as the monitor is enabled.
tBGTask	The background task (tBGTask) is the core task switching process. It receives packets from the hardware ASICs and switches them to the appropriate functional task to process that packet type or group. The tBGTask typically runs with a high CPU utilization (90% or greater). It is constantly checking for packets to be sent up by the hardware ASICs. It only releases control of the CPU if packets are sent to the switch or if timer functions signal another task to become active.
tCardTask	The I/O card event task (tCardTask) manages the event signaling hardware and state machine for the I/O cards in a chassis-based system.
tChecksumPoll	The "*i*" series chipset checksum polling task (tChecksumPoll) periodically polls the "*i*" series chipset boards for fabric checksum errors.
tConsole	The console task.
tdiagTask	The diagnostic task (tdiagTask) executes the diagnostic routines for the particular hardware platform.
tDvmrpTask	The distance vector multicast routing protocol task (tDvmrpTask) implements and processes DVMRP on the switch.
tEapsTask	The Ethernet automatic protection switching task implements and processes EAPS on the switch.
tEdpTask	The Extreme Discovery Protocol task (tEdpTask) implements and processes the EDP neighbor discovery process.
tEsrpTask	The Extreme Standby Router Protocol (tEsrpTask) implements and processes ESRP on the switch.
tExcTask	If the operating system recognizes an exception condition, it will invoke the exception handling task (tExcTask).
tFastTimer	The fast timer task (tFastTimer) is used to maintain a queue of timer events triggering periodic or single event functions. These events have a small delay in time between re-occurrences. The tFastTimer has a higher priority than the slow timer task (tSlowTimer). Therefore, tFastTimer events are processed prior to iSlowTimer events occuring at the same time.
tfdbAgeTask	The forwarding database aging task (tfdbAgeTask) performs the aging of MAC FDB entries in the hardware and software tables.
tIpxTask	The IPX input task (tIpxTask) handles inbound IPX control packets such as RIP, SAP, and Xping.

Table 26: ExtremeWare Task Descriptions (continued)

Task	Description
tIpxTx	The IPX transmit task (tIpxTx) handles the IPX transmission of control packets such as RIP and SAP.
tIquery	The iQuery support task for 3DNS (tIquery) processes iQuery requests.
TIRDP	The ICMP router discovery protocol task (tIRDP) implements and processes IRDP on the switch.
tISRtask	The interrupt service routine task (tISRtask) manages the interrupt driven port link state changes.
tLinkEvent	The link event task (tLinkEvent) is the interrupt driven link event processing task. It handles hardware interrupts for link events.
tMACPoll	The media access controller poll task (tMACPoll) polls the various MAC PHY chips on the switch to pull up MAC Layer control messages for the CPU to process.
tmt32LinkPoll	F32F card link poll task.
tmuTelnetd	The telnet daemon task.
tNetTask	The network stack task (tNetTask) handles all the software-based processing of packets including: • Packets that cannot be handled by the switch's ASIC because the forwarding tables do not have entries built in. • Packets destined to the CPU for one of the router interfaces. • Packets that must be examined or snooped by the CPUPackets detected for copying to the CPU.
tNMCEvent	The network management controller event task (tNMCEvent) manages event signaling hardware and state machine on a BlackDiamond switch’s redundant MSM CPU cards.
tOpenPort	A server load balancing (SLB) Layer 4/Layer 7 health check sub-task.
tospfMsgTask	The OSPF message processing task (tospfMsgTask) implements and manages the processing of OSPF messages.
tospfSpfTask	The OSPF shortest path forward task (tospfSpfTask) executes the SPF algorithm run processing for OSPF.
tospfTimer	The OSPF timer task (tospfTimer) manages the internal timer trigger functions and delays for OSPF.
tPCSPoll	The tPCSPoll task services the Gigabit Ethernet PCS poll messages.
tPhyPoll	The PHY layer poll task (tPhyPoll) polls the Road Runner PHY layer every 2 seconds to verify the proper operation.
tPimTask	The protocol independent multicast task (tPimTask) implements and processes PIM on the switch.
tPingServer	The server load balancing (SLB) Layer 3 ping health check sub-task.
tPortProbe	A server load balancing (SLB) Layer 4/Layer 7 health check sub-task.
tPortUtilization	The port utilization data collection task (tPortUtilization) is a 30 second task that pulls physical port data statistics from the hardware and updates the software database tables.
tRip	The Routing Information Protocol task (tRip) implements and processes RIP on the switch.
tRipTimer	The RIP timer task (tRipTimer) manages the internal timer trigger functions and delays for RIP.
TRmonTask	The remote monitoring task
tRRPoll	The Road Runner poll task (tRRPoll) pulls the MAC and PHY layer statistics from the store in the software based tables.

Table 26: ExtremeWare Task Descriptions (continued)

Task	Description
tRxMsgTask	The receive message task (tRxMsgTask) is located on the secondary system. ExtremeWare 6.2 commences use of the secondary CPU in BlackDiamond switches. This is the secondary slave CPU inter-CPU receive task.
tServAlive	The server load balancing (SLB) health check server task.
tShell	The core operating system internal shell process (tShell) is spawned whenever the internal shell is accessed.
tSlbFailover	The server load balancing failover task.
tSlowTimer	The slow timer task (tSlowTimer) maintains a queue of timer events triggering periodic or single event functions. Typically these events have a large period gap in terms of time between reccurrences.
tsmartTrap	Extreme smart trap task.
tSnmpd	The SNMP daemon task manages all SNMP processing on the system.
tSntpc	The simple network time protocol client task (tSntpc) implements the SNTP client function and processing.
tsshshell	The secure shell (SSH) task.
tStatsPoll	The port interface statistics poll task (tStatsPoll) polls the port interfaces for statistic counters.
tstpTask	The Spanning Tree protocol task (tstpTask) implements the STP algorithm and processing.
tSwFault	The software fault handler task (tSwFault) will perform a stack dump for any task that has crashed.
tsyslogTask	The system log task (tsyslogTask) receives messages/text from other tasks and asynchronously logs these to the switch NVRAM log area.
tTimeout	The Timeout task (tTimeout) is used to manage and execute various functions on timeouts.
tTRRecv	The trace route receiver task (tTrRecv) is spawned dynamically when the trace route utility is used.
tvrrpTask	The virtual router redundancy protocol task (tvrrpTask) implements and processes VRRP on the switch.

Investigate tasks that, for no apparent reason, show CPU utilization consistently above 25% (except for the BGTask). Configure the appropriate debug-trace command and look for messages indicating a problem. Common problems are source or destination addresses.

Example

The following command displays the show command output on the switch:

```
top
```

The output of this command looks similar to the following:

```
Total number of tasks: 46
Task Name       Task Id         Task PC         Status   % CPU Max % util
================================================================================
   tBGTask      836f18e0        80748f98        READY        99   99
  tExcTask      8137ce90        8075ab2c        PEND          0    0
  tLogTask      8135e2a0        8075ab2c        PEND          0    0
tSlowTimer      813ccf50        8075ab2c        PEND          0    0
tFastTimer      813ff1f0        8075ab2c        PEND          0    0
```

```
  tTimeout      81384f50        8075ab2c        PEND          0    0
tsyslogTas      81389660        8075ab2c        PEND+T        0    0
tledPollTa      81390ef0        8075ab2c        PEND          0    0
tAsyncSave      814feb10        8075ab2c        PEND          0    0
 tpifstate      81a85590        8075ab2c        PEND          0    0
  tbgpTask      81eaacd0        807169f4        PEND+T        0    0
tbgpTimerT      81eaecd0        80749164        DELAY         0    0
 tBgQosMon      81eb6be0        8075ab2c        PEND          0    0
 tEapsTask      82bd2a00        8075ab2c        PEND          0    0
  tSwFault      82c75530        8075ab2c        PEND          0    0
tFdbAgeTas      82c85530        8075ab2c        PEND          0    0
tFdbSyncTa      82c89530        807489a0        SUSPEND       0    0
 tdiagTask      82c8d620        8075ab2c        PEND          0    0
  tIpxTask      82c91620        8075ab2c        PEND          0    0
    tIpxTx      836e97f0        8075ab2c        PEND          0    1
Press ‘h’ for help
```

History

This command was available in ExtremeWare 4.0.

Platform Availability

This command is available on all “*i*” series platforms.

Index of Commands

D

E

T

U

X